Can Bitcoin Survive the Age of Quantum Computing?

The Quantum Threat and the Evolution of Bitcoin Security

About

Metanomia is a think tank focused on analyzing transformations in the cryptocurrency market and related technologies.

Researchers

Yonkyung Kim

Hyemin Son

Corresponding Author

Taemin Oh

Contents

Introduction

Has Bitcoin ever been hacked? Looking only at news headlines, it is easy to believe so. In 2013, as Bitcoin's price fluctuated following the Mt. Gox security incident amid market overheating, CNBC ran a headline that plainly stated, "Bitcoin Hacked".[1] Reading just the title, one might assume the Bitcoin protocol itself had been breached. If the 2013 incident was a one-off security failure, the 2014 Mt. Gox crisis was closer to a total collapse. At the time, Mt. Gox was the world's largest Bitcoin exchange, and approximately 850,000 BTC—including customer assets—vanished.[2] This amounted to around $480 million at the time, representing roughly 7% of the total circulating supply. While many reports accurately described it as an "exchange bankruptcy" or "exchange hack," the public simply internalized it as a "Bitcoin hack". This narrative repeated itself in 2016 when another major exchange, Bitfinex, saw roughly 120,000 BTC stolen.[3]

However, what was breached in these incidents was not the Bitcoin network, but the servers of centralized exchanges. Exchanges are private companies that pool and store users' bitcoins on centralized servers. To use an analogy, a depository holding gold belonging to multiple people was robbed; the fundamental properties of gold itself were not compromised. Hackers exploited vulnerabilities in corporate security systems—specifically server management and internal controls—not Bitcoin's cryptographic architecture.

The Bitcoin network itself has not been hacked a single time since its launch on January 3, 2009.[4] It is worth reflecting on what this means. For over 16 years, the Bitcoin network has generated a block every 10 minutes without fail. The value of the assets secured on this network has reached a market capitalization of approximately $1.6 trillion (as of May 2026). For anyone capable of breaking this network's cryptography, the potential reward would be immense. Despite huge incentives for countless attackers—ranging from independent hackers and research institutions to state-level actors—the system has never been breached.

But does this robustness have an expiration date? This is a valid question because Bitcoin's defenses rest entirely on a single premise: that the mathematics underpinning its security remain unbroken. If that math is broken, a 16-year flawless track record offers no guarantees.

This report examines that premise in light of the emergence of quantum computing. We will sequentially explore the threat quantum computers pose to Bitcoin's cryptographic systems, how close this threat is to reality, and how the Bitcoin community is preparing for it.

To state the conclusion upfront: quantum computing is not exclusively a Bitcoin problem, but a challenge for our entire modern digital civilization. Furthermore, Bitcoin already possesses the structural pathways necessary to address this challenge.

01

Bitcoin's Two Pillars of Security

Bitcoin is a decentralized electronic cash system where transactions occur without a central authority. Even though there is no single entity—like a bank or a registry—to verify identities and manage the ledger, the network enforces that specific Unspent Transaction Outputs (UTXOs)[1] only move when valid conditions are met. Here, "valid" does not mean the owner's identity has been verified; rather, it means cryptographic proof has been provided showing control over the private key capable of unlocking the asset's locking condition.

Bitcoin's security relies on two mathematical structures. The first is Elliptic Curve Cryptography (ECC), which protects the control over assets—effectively the economic authority corresponding to ownership—through a system where digital signatures generated by private keys are verified by public keys. The second is the SHA-256 hash function, which is used in mining and linking blocks together, thereby supporting the cost structure of consensus and the integrity of the records. Combining these two pillars, Bitcoin simultaneously achieves proof of ownership without banks and a virtually immutable record.

Keys and Locks

In traditional finance, ownership relies on third-party verification and record-keeping. Bank ledgers for deposits, registries for real estate, and depository or account records for stocks serve as the ultimate arbiters. The authority to determine "who the owner is" is concentrated in specific institutions. In contrast, Bitcoin transactions are executed without institutional identity verification, seals, or paperwork. This is because control over a Bitcoin asset is established through cryptographic proof rather than human identity. What the network verifies is not "who this person is," but "can they produce a valid signature with the private key?"

Bitcoin's ownership system is based on the mathematical principles of public-key cryptography. A user first determines a private key and mathematically derives a public key from it. This relationship is often compared to a lock and key. The private key is the key accessible only to the owner. This key is a 256-bit number, and the number of possible combinations is approximately 2^{256}, or roughly 1.15×10^{77}. This figure is large enough to rival estimates of the total number of atoms in the observable universe. Such an astronomical number of possibilities means that guessing someone else's private key through brute force is practically impossible. The public key is mathematically derived from this private key and acts as the lock. Anyone can see this lock, but seeing it

1) UTXO (Unspent Transaction Output) refers to a transaction output that has not yet been spent. Rather than maintaining account balances, Bitcoin operates on a transaction-based model in which outputs from previous transactions are used as inputs to new transactions. A user's bitcoin holdings are therefore represented by the collection of unspent transaction outputs associated with them.

alone is not enough to unlock it.

To transfer bitcoin, one must generate a digital signature using the private key. A digital signature is a piece of data that mathematically proves the "authority to satisfy the locking condition"—in other words, the fact that "I control the private key and I authorize this transaction". Network participants use the public key to verify that the signature is correct. This mechanism works because there is an asymmetry between the private key and the public key. Calculating a public key from a private key is simple, but reverse-engineering a private key from a public key is practically impossible with current computers. This very asymmetry is the core premise underpinning the security of Bitcoin's digital signatures.[5]

The First Pillar: Elliptic Curve Cryptography

The property that it is easy to calculate a public key from a private key, but practically impossible to do the reverse, is specifically implemented on a mathematical framework known as "elliptic curve cryptography". Since its inception, Bitcoin has used the Elliptic Curve Digital Signature Algorithm (ECDSA) operating on a specific elliptic curve called secp256k1.[6] The secp256k1 elliptic curve has a fixed point known as the generator point. In Bitcoin, a public key is created by applying a private key, which is a number, to this generator point. Mathematically, this is the result of repeatedly applying the elliptic curve's addition rules to the generator point as many times as the value of the private key, and the resulting point on the curve becomes the public key. This computation is performed very quickly. However, reverse-engineering how many operations were applied—that is, figuring out what the private key is—just by looking at the resulting public key is practically impossible for today's classical computers. This is known as the Elliptic Curve Discrete Logarithm Problem (ECDLP). As long as the computational difficulty of the ECDLP holds, it remains practically impossible for an attacker to recover the private key or forge a valid signature, even if the public key is exposed. This is the first mathematical pillar supporting Bitcoin's protection of control.

[Table 1] **How ECC is Used in Bitcoin**

Application Area	Underlying Technology	Function
Public Key Generation	secp256k1 Elliptic Curve Operations	Derives a public key from a private key
Transaction Signing	ECDSA / Schnorr	Proves control and authorizes transactions

The Second Pillar: SHA-256

The second pillar of Bitcoin's security is the SHA-256 hash function. A hash function is a one-way function that takes an input of arbitrary length and outputs a hash value of fixed length. The same input will always produce the same output, but even the slightest change in the input generates a completely different hash value. For instance, "Hello" and "Hello." differ by only a single

period, but the resulting hash values are entirely different. Because of this property, hashes are widely used to verify whether original data has been tampered with.

In Bitcoin, SHA-256 serves two core roles. First, it acts as the foundation of Proof of Work (PoW). In Bitcoin, the entity that will generate the next block is determined not by specific approval or authority, but through a computational competition to find a hash value that meets a certain difficulty condition. A miner bundles transaction records and various information in the block, then repeatedly alters adjustable values within it—most notably an arbitrary number called a nonce—to find a hash value satisfying the difficulty conditions demanded by the network. This process is closer to a brute-force search that attempts repeatedly until a matching value emerges, rather than a logical deduction to find an answer. Therefore, practically the only way to increase the probability of success is to perform more hash calculations. Thanks to this structure, block generation incurs realistic electricity and equipment costs, and which block is accepted as legitimate is ultimately governed by the actual computational cost expended.

Second, SHA-256 supports the integrity of records. In a blockchain, each block is sequentially linked by including the hash of the previous block. Consequently, if even a single bit of a past block's content is altered, the hash of that block changes, which subsequently breaks the connection to all following blocks that reference that value. Tampering with a single past block practically means having to recreate all the blocks that follow it. However, in Bitcoin, this is not merely a matter of reorganizing data. An attacker would have to redo the Proof of Work for the altered blocks while simultaneously catching up to the accumulated computational power of the honest chain, which continues to extend in the meantime. Ultimately, overturning historical records requires the expenditure of astronomical computational resources and time. This is why Bitcoin's records are virtually immutable.

It is precisely in this regard that SHA-256 serves as the second mathematical pillar of Bitcoin security. If elliptic curve cryptography determines the legitimate control required to move assets, SHA-256 dictates what records will remain as legitimate history. One upholds control, while the other underpins the continuity and cost structure of the record. Bitcoin's security holds only when these two pillars stand together.

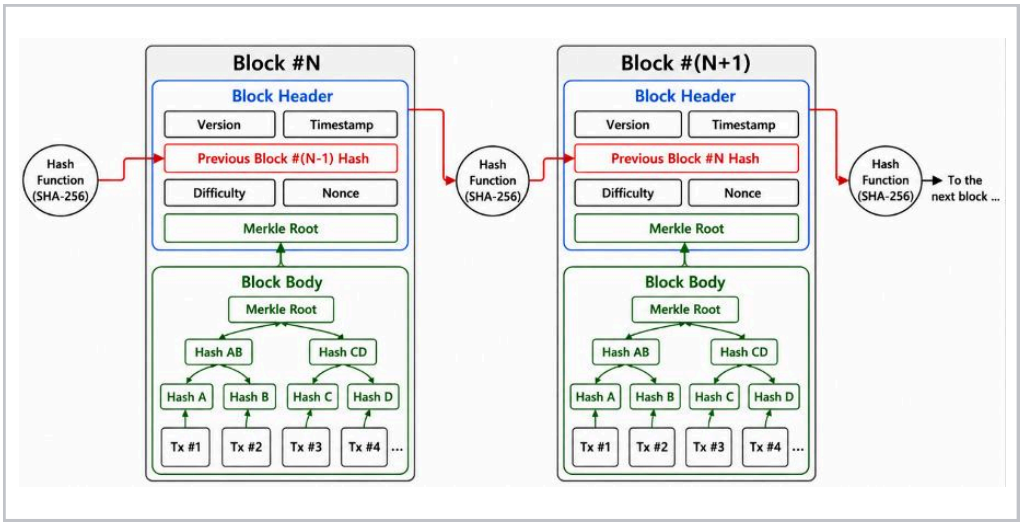

[Figure 1] **Structure of Block Connection via Previous Block Hash**

Each block's header serves as the input to a SHA-256-based hash function, and the resulting value is stored in the "Previous Block Hash" field of the next block's header. The transaction records in the block body are summarized into a single Merkle Root via a Merkle Tree, and this value is also included in the block header. Therefore, if any transaction record or block header field is altered, the hash of that block changes, and the chain of connections to all subsequent blocks is affected in turn.

02

What Quantum Computers Target

2)

Quantum computing is a computational model that performs calculations using the principles of quantum mechanics. In the early 1980s, Richard Feynman proposed the need for a computational device capable of efficiently simulating quantum systems. Subsequently, in 1985, David Deutsch formalized the theoretical model of a universal quantum computer. Since then, the field has evolved into a distinct computational paradigm through the development of various quantum algorithms.

The two mathematical pillars of Bitcoin discussed earlier rely on the premise that it is practically impossible to reverse-engineer a private key from a public key, or to accelerate hash searches enough to neutralize the cost structure of Proof of Work. However, quantum computing*2)* research, which emerged in the 1980s, revealed the potential to fundamentally shake this premise with the introduction of Shor's Algorithm in the mid-1990s.[7] As it became clear that quantum computers could solve specific mathematical problems at speeds incomparable to classical computers, the entirety of modern cryptography, including Bitcoin, was placed in a new threat environment.

The Difference Between Classical and Quantum Computers

The computers we use daily process information in units called bits. A bit can only exist in one of two states, 0 or 1, and all calculations are composed of combinations of these binary states.

In contrast, the fundamental unit of a quantum computer is the "qubit." Until it is measured, a qubit does not settle into a definitive 0 or 1, but can exist as a combination of both states. This is called "superposition." If a classical bit is like a coin lying heads or tails, a qubit is more akin to a coin spinning in the air. However, while a spinning coin's outcome is merely hidden, a qubit's superposition is fundamentally different in that it is a quantum state described as a combination of two states prior to measurement.

Quantum systems also exhibit a phenomenon known as "entanglement." Entanglement occurs when two or more qubits form a single, combined quantum state, making it impossible to describe the state of one qubit independently of the others. In this state, measuring one qubit reveals a strong correlation with the measurement of the other, a correlation that persists even if the qubits are physically separated. If superposition is a concept that explains the state of an individual qubit, entanglement explains the phenomenon where multiple qubits form a single, inseparable combined state.

When superposition and entanglement are combined, quantum computers perform calculations in a fundamentally different way than classical computers. Parallel computing in a classical computer involves multiple processors dividing and handling different tasks simultaneously. A quantum computer, on the other hand, uses superposition and entanglement to configure its qubits into a single

quantum state encompassing the possibilities of multiple computational paths, and then applies a quantum operation to that entire state. During the subsequent calculation process, "interference" occurs, where different quantum states reinforce or cancel each other out. Quantum algorithms exploit this to amplify the probability of reaching the correct answer and reduce the probability of unwanted outcomes.

Thanks to this architecture, quantum computers can arrive at solutions for specific problems in significantly fewer computational steps than classical computers. Of course, quantum computers are not faster at every problem. Their overwhelming computational advantage is limited to specific types of mathematical problems. The critical point is that several problems foundational to modern cryptography fall precisely into this category. Notably, prime factorization and the discrete logarithm problem are known to be efficiently solvable by quantum algorithms.

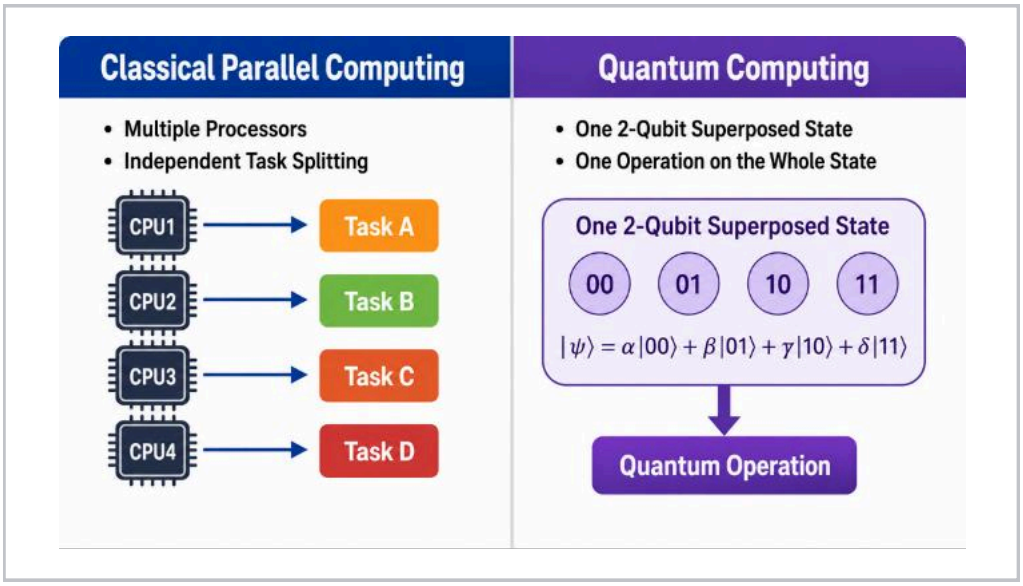

[Figure 2] **Classical Computer's Parallel Processing and Quantum Computer's Superposition State Calculation**

Unlike classical parallel computing, which uses multiple processors to divide and handle different tasks, a quantum computer applies a quantum operation to a single superposition state containing multiple possibilities. During the subsequent calculation process, interference between quantum states is used to increase the probability of measuring the desired result and decrease the probability of unwanted results. This figure conceptually illustrates the difference in these computational methods.

Shor's Algorithm and Grover's Algorithm

The threat quantum computers pose to Bitcoin splits into two main vectors. One is Shor's Algorithm, which targets elliptic curve cryptography, and the other is Grover's Algorithm, which accelerates hash searches. Both algorithms impact Bitcoin's cryptographic foundations, but the nature of their threats differs. While Shor's Algorithm directly targets Bitcoin's control structure, Grover's Algorithm primarily weakens the cost structure of Proof of Work and hash searching.

Proposed by Peter Shor in 1994, Shor's Algorithm demonstrated that if a

sufficiently large, fault-tolerant (error-corrected) quantum computer exists, it can solve discrete logarithm and integer factorization problems far more efficiently than a classical computer.[8] Bitcoin's digital signature scheme is based on the discrete logarithm problem over the secp256k1 elliptic curve. In other words, it relies on the premise that deriving a private key from a public key is hard.

If a quantum computer capable of running Shor's Algorithm emerges, this premise would no longer hold, as calculating a private key from a public key becomes theoretically possible. This can be likened to a technology that can forge a key simply by looking at the lock. In Bitcoin, the authority to move an asset is proven by the fact that one knows the private key. If calculating the private key from the public key becomes possible, assets in addresses where the public key has already been exposed are at risk of being moved by an attacker acting as the legitimate owner. In this regard, Shor's Algorithm directly targets Bitcoin's first pillar—the control protection structure based on elliptic curve cryptography.

Conversely, Grover's Algorithm, proposed by Lov Grover in 1996, poses a different kind of threat to hash functions. Grover's Algorithm is a quantum algorithm that reduces the number of attempts required for a brute-force search to roughly its square root.[9] When trying to find a value that satisfies a specific condition, if a classical computer requires N attempts, Grover's Algorithm reduces the required attempts to approximately $\sqrt{N}$.

Applying this to SHA-256 effectively lowers its 256-bit security strength to around a 128-bit level. This is certainly a significant weakening. However, 128-bit security remains exceptionally high. 2^{128} is an astronomically large number (approximately 3.4×10^{38}), and it is practically unsearchable with existing computational resources. The quantum threat to SHA-256 is more accurately described as a reduction in computational difficulty rather than an immediate collapse of the hash function.

Ultimately, the threat quantum computers pose to Bitcoin does not apply equally to its two pillars. SHA-256 merely sees its search difficulty lowered by Grover's Algorithm; it is not a structure that collapses instantly. In contrast, elliptic curve cryptography is the direct target of Shor's Algorithm and faces a far more fundamental and direct threat: the ability to derive private keys from public keys if a sufficiently powerful quantum computer emerges.

Therefore, the core of the quantum threat lies in the vulnerability of control over addresses whose public keys have been exposed. However, this is strictly a theoretical possibility for now. For such an attack to actually occur, a large-scale, fault-tolerant quantum computer capable of running Shor's Algorithm is required. The question we must now explore is: exactly how close is this threat to becoming reality?

[Table 2] **Comparison of Shor's Algorithm and Grover's Algorithm**

Category	Shor's Algorithm	Grover's Algorithm
Target	Elliptic Curve Cryptography (ECC)	Hash Function (SHA-256)
Effect	Public key → Private key reverse engineering (Fundamental break)	Security strength reduced from 256-bit to 128-bit (Partial weakening)
Analogy	Forging a key by looking at the lock	Speeding up attempts to crack a safe's passcode
Threat Urgency	High	Relatively Low

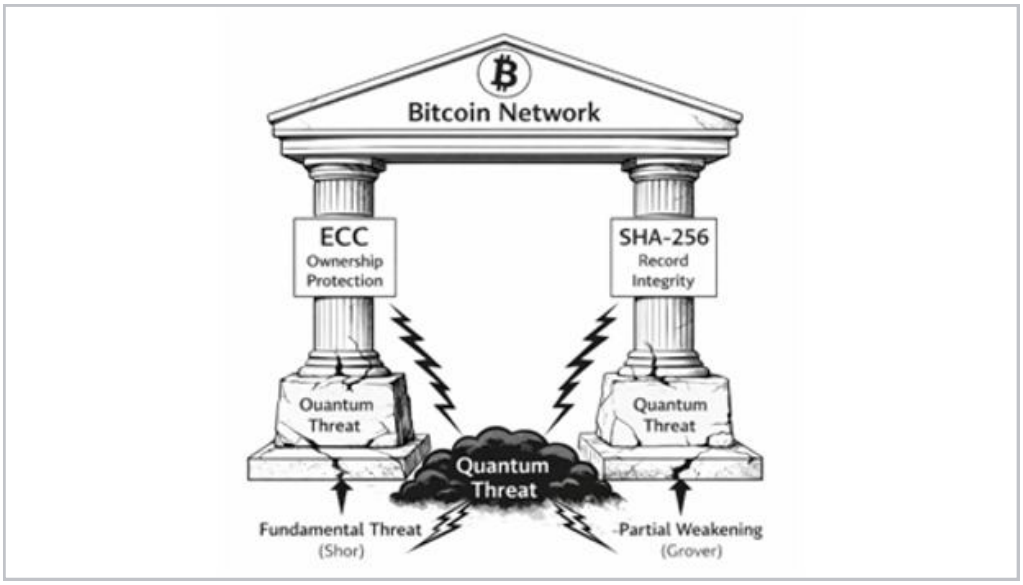

[Figure 3] **Quantum Threats to Bitcoin's Cryptographic Foundation**

03

The Quantum Threat: A Defense Already in Progress

The State of Quantum Computing: Advances and Limitations

The prospect that quantum computers could pose a long-term threat to Bitcoin's cryptographic system is no longer confined to the realm of science fiction. In recent years, major technology companies have successively announced achievements demonstrating the scale and error-correction potential of quantum computers. In 2023, IBM unveiled Condor, a 1,121-qubit quantum processor, breaking the 1,000-qubit barrier.[10] Subsequently, in its updated 2025 roadmap, IBM set a goal to develop the Starling system by 2029, which aims to perform 100 million quantum gate operations using 200 logical qubits.[11] Google announced Willow, a quantum chip equipped with 105 qubits, in December 2024.[12] Microsoft also announced the Majorana 1 chip in February 2025, stating it had laid the groundwork for scaling up to 1 million qubits on a single chip in the long term.[13]

[Table 3] **Comparison of Major Quantum Computing Architectures**

Category	Condor	Willow	Majorana 1
Developer	IBM	Google	Microsoft
Release Year	2023	2024	2025
Qubit Scale	1,121 qubits	105 qubits	8-qubit experimental chip; long-term target of 1 million qubits
Technology Type	Superconducting qubits	Superconducting qubits	Topological qubits
Core Significance	Large processor breaking the 1,000-qubit barrier	Advances in error correction and logical qubit stability	Introduces a new qubit architecture with enhanced error resistance
Current Limitations	High qubit count, but not yet fault-tolerant	Shows strong benchmark results, which do not translate to real-world decryption capability	Currently at 8 qubits; large-scale expansion remains a long-term challenge

Among these, Google's Willow serves as a prime example illustrating the significance of recent advancements in quantum computing. According to Google, Willow performed a calculation in just 5 minutes that would take an estimated 10^{25} years on a classical computer. However, Willow's significance lies not merely in computational speed; a more crucial achievement was demonstrating the so-called "below threshold" phenomenon in error correction experiments, where increasing the number of qubits actually reduced the error rate. This demonstrated that adding more qubits does not simply accumulate errors; rather, within an appropriate error-correction framework, the stability of logical qubits can actually improve.

Nevertheless, this achievement does not immediately translate to decryption capabilities, as the problem Willow solved was a benchmark designed specifically to measure quantum computer performance. Google also made this clear in interviews following the announcement, with a spokesperson stating, "Current estimates suggest it will take at least another 10 years to break RSA*3)*, and this would likely require about 4 million physical qubits*4)*."[14] Thus, while Willow's performance is a major step forward showcasing the potential of quantum computers, it is still far from the practical ability to break cryptographic systems.

In the case of Bitcoin, the scale of a quantum computer required for decryption is massive; researchers estimate that millions to tens of millions of physical qubits would be needed to attack the elliptic curve signature scheme Bitcoin uses. A research team at the University of Sussex, basing their resource estimates on specific error-correction methods and hardware assumptions, estimated that approximately 13 million physical qubits could be required to attack Bitcoin's secp256k1-based elliptic curve signature scheme within a single day.[15] Considering that Google's Willow chip is at a 105-qubit scale, the gap to the required level is roughly 120,000-fold. This gap illustrates that a significant distance remains between current quantum computing technology and practical cryptographic attacks.

To understand why the number of qubits required for decryption is so large, one must consider the error rate problem inherent in quantum computers. Because qubits rely on highly sensitive quantum states, they are easily prone to errors even from minor environmental interference. Consequently, performing actual calculations requires a large number of additional physical qubits dedicated to error correction. For this reason, quantum computing research distinguishes between "physical qubits," which are actually implemented in hardware, and "logical qubits," which are stable computational units constructed through error correction. Depending on the error-correction scheme and the underlying error rate, implementing a single stable logical qubit can require hundreds to thousands of physical qubits. The "13 million qubits" estimate provided by the University of Sussex team also refers to the scale of physical qubits, including those needed for error correction, not logical qubits. Furthermore, translating this into a physical device requires not only the qubits themselves but also a massive hardware infrastructure for control, measurement, cooling, and wiring.

Due to these technical limitations, current quantum computers are often described as being in the "NISQ (Noisy Intermediate-Scale Quantum)" era. This term, proposed by physicist John Preskill, refers to early-stage, noisy quantum

3)

RSA is a widely used public-key cryptosystem proposed in 1977 by Ronald Rivest, Adi Shamir, and Leonard Adleman. Its security is based on the computational difficulty of factoring a large composite number n—formed by multiplying two large prime numbers, p and q—back into its original prime factors. However, a sufficiently large fault-tolerant quantum computer could solve this problem efficiently using Shor's Algorithm. Although RSA is not the digital signature scheme used by Bitcoin, it is frequently cited as a representative example when discussing the cryptographic threats posed by quantum computing.

4)

A physical qubit is an individual qubit implemented in actual hardware. In contrast, a logical qubit is a stable computational unit formed by combining multiple physical qubits through quantum error-correction techniques, enabling the detection and correction of errors that arise during computation.

computers possessing anywhere from 50 to a few hundred qubits.[16] While NISQ devices can physically realize phenomena like quantum superposition and entanglement, their high levels of noise and errors make it difficult to perform stable calculations over extended periods. Thus, while current quantum computers show meaningful results in specific experimental problems or limited optimization calculations, they are still technically ill-equipped to handle massive computations like breaking Bitcoin's cryptographic system.

However, this does not mean the quantum threat can be ignored; the critical issue is not whether quantum computers can break Bitcoin immediately, but the difficulty of accurately predicting when that threat will materialize. Quantum computing research is advancing rapidly, and issues surrounding error correction and hardware scalability are gradually improving. Ultimately, the quantum threat is a matter of a timeline, not an imminent collapse. For Bitcoin to endure as a global financial infrastructure in the long run, it must prepare for a post-quantum transition in advance. This transition is not merely about swapping out a single cryptographic algorithm; it is a long-term migration process requiring wallets, nodes, exchanges, users, and protocol consensus to move in unison. If preparations are made early, the transition will be an orderly upgrade, but if delayed, it becomes a race against time to outmaneuver attackers equipped with quantum computers.

The Quantum Threat Goes Beyond Bitcoin

An important point must be made here: the cryptographic threat posed by quantum computers is not confined to cryptocurrency. Bitcoin utilizes the secp256k1-based elliptic curve signature scheme and the SHA-256 hash function, but these are not isolated technologies created specifically for Bitcoin. More accurately, Bitcoin is built upon the core technological lineages of modern digital security infrastructure: public-key cryptography, digital signatures, and cryptographic hash functions.

SHA-256 is a representative cryptographic hash function belonging to the SHA-2 family, an algorithm designed by the National Security Agency (NSA) and published as a standard by the National Institute of Standards and Technology (NIST). Today, SHA-2 family hash functions are used in diverse areas such as digital signatures, certificates, software integrity verification, and secure communications, serving as a cryptographic foundation for data protection in online payments and cloud services. Transport Layer Security (TLS), the core protocol for modern internet security, also operates on top of these cryptographic technologies. TLS is the standard protocol for encrypting internet communications and forms the security foundation for HTTPS traffic. When a user accesses a website starting with HTTPS in their browser, TLS verifies the server's identity through its certificate and agrees on a session key for secure communication. This process utilizes public-key cryptography, digital signatures, certificates, and hash functions in tandem; furthermore, TLS is not a technology used solely for website access. The entire internet infrastructure—including email security, cloud services, Application Programming Interface (API) communications, and corporate networks—relies on TLS and its surrounding authentication frameworks.

At the center of this process is the Public Key Infrastructure (PKI). PKI is a system that uses digital certificates to guarantee that a specific public key actually belongs to a particular server or organization. Web browsers verify server certificates based on root certificates from trusted Certificate Authorities, thereby forming the chain of trust for internet communications. This structure has long relied on RSA or elliptic curve-based public-key cryptography and digital signatures. In other words, behind everyday actions like accessing websites, installing apps, receiving software updates, and storing data in the cloud, there is an invisible architecture of trust constructed by public-key cryptography and hash functions.

The traditional financial system is also not exempt from this structure; online banking authentication, the protection of financial transaction messages, card payment networks, international settlement systems, electronic document authentication, and software code signing all rely deeply on public-key authentication, digital signatures, and encrypted communication technologies. The modern financial system operates atop a cryptographic trust architecture combining servers, certificates, keys, signatures, hashes, and security protocols. Therefore, if a sufficiently powerful quantum computer emerges, Bitcoin will not be the only entity affected. A vast array of fields that constitute digital trust—including secure internet communications, digital signatures, certificates, financial transactions, software distribution, and government and public sector security systems—would be simultaneously impacted. What is vulnerable to quantum computers is the entirety of modern civilization's digital security framework, and paradoxically, this very fact provides a degree of buffering for Bitcoin. This is because the Bitcoin community is not alone in needing to solve this problem; governments, enterprises, and research institutions worldwide are concurrently tackling the exact same issue.

The Transition to Post-Quantum Cryptography Has Already Begun

This collective response is already underway; in August 2024, after an eight-year international standardization process, NIST finalized and announced the first three Post-Quantum Cryptography (PQC) standards.[17]

The first is ML-KEM (Module-Lattice-Based Key-Encapsulation Mechanism), a key encapsulation method designed to allow communicating parties to securely share secret keys, intended to replace or supplement key exchange structures used in existing internet security protocols for a post-quantum environment. The second is ML-DSA (Module-Lattice-Based Digital Signature Algorithm), a lattice-based digital signature scheme aimed at replacing existing RSA or ECDSA signatures. The third is SLH-DSA (Stateless Hash-Based Digital Signature Algorithm), a hash-based digital signature scheme built on a different mathematical approach than ML-DSA, proposed as a supplementary standard in case ML-DSA proves vulnerable.

[Table 4] **Three Post-Quantum Cryptography Standards Announced by NIST**

Standard Name	Formal Name	Purpose	Mathematical Basis	Description
FIPS 203	ML-KEM	Key Establishment	Lattice Problems	Secure sharing of secret keys
FIPS 204	ML-DSA	Digital Signatures	Lattice Problems	Verification of authorship
FIPS 205	SLH-DSA	Digital Signatures	Hash-Based	Long-term backup security measure

Post-Quantum Cryptography does not refer to a single algorithm; rather, it is a collection of cryptographic techniques based on various mathematical problems that are difficult for quantum computers to solve efficiently. Currently, one of the most prominent approaches is lattice-based cryptography. Lattice-based cryptography derives its security from the computational difficulty of high-dimensional lattice problems and is considered a strong candidate for both key establishment and digital signatures. However, it carries the burden of increased key and signature sizes, which can raise storage and communication costs during system implementation. Hash-based signatures offer the advantage of relatively simple security assumptions and well-studied stability over long periods, but they can be disadvantageous regarding signature size and computational efficiency. For this reason, future security infrastructure is highly likely to evolve by combining multiple cryptographic techniques with different mathematical foundations rather than relying on a single technology.

The important point is that this transition is no longer confined to the research phase. NIST has stated that these new standards are ready for immediate use and has urged computer system administrators to begin transitioning to the new standards as soon as possible. Major tech companies have also begun implementing post-quantum cryptography on an experimental or partial basis. Google and Cloudflare have tested and partially deployed post-quantum key exchange methods in internet communication protocols like TLS,[18] and Apple introduced the PQ3 protocol to iMessage, which incorporates post-quantum cryptography.[19] IBM is also advancing post-quantum security infrastructure, including a quantum-safe TLS mode for IBM Cloud and a PQC transition plan for the IBM Quantum Platform.[20] Furthermore, Zoom announced that it is offering post-quantum end-to-end encryption (E2EE) features in Zoom Workplace and Zoom Meetings.[21]

This trend reveals an important truth: for now, the spread of quantum-resistant cryptography is outpacing the development of quantum computers capable of real-world attacks. In effect, the shield is being institutionalized before the spear. Long before quantum computers become an actual threat, the global internet infrastructure and financial systems have already begun migrating to new cryptographic regimes.

This international transition holds significant implications for Bitcoin; if Bitcoin must transition to post-quantum cryptography in the future, it is highly unlikely to be forging an entirely new path alone, but rather following a path the world is already taking. However, Bitcoin's transition is not a system update that can be centrally mandated by a government or corporation. It is a matter of consensus

—one that requires wallets, nodes, exchanges, miners, users, and the developer community[22] to all move in the same direction. Therefore, the remaining challenge for Bitcoin lies in determining the method, timing, and social consensus through which this post-quantum transition will be achieved.

04

Bitcoin's Transition Path

If quantum computers pose a realistic threat, Bitcoin's long-term countermeasure entails migrating from its current elliptic curve-based signature scheme to a post-quantum cryptography (PQC) signature scheme. However, the core challenge lies not in the theoretical possibility of the transition itself, but rather in the methodology of implementing it within Bitcoin's decentralized consensus architecture. This is not a straightforward software update; it is a long-term evolutionary process that requires the synchronized alignment of protocol rules, wallet and exchange infrastructures, user behavior, and broad community consensus.

Bitcoin's History of Upgrades

Bitcoin is commonly perceived as a system immutable by design, built upon unalterable rules. This conventional wisdom is valid to a certain extent. The foundational components that govern the network's core order—such as the hard-capped money supply, the Proof-of-Work (PoW) mechanism, difficulty adjustment cycles, and transaction verification rules—are notoriously resistant to change. However, this does not mean the Bitcoin protocol is entirely static. More precisely, Bitcoin is a system that "cannot be changed arbitrarily." Alterations are entirely possible, but the process mandates public proposals, rigorous technical scrutiny, code implementation, voluntary adoption by network participants, and a comprehensive social consensus.

These protocol modifications are typically debated through a standardized framework known as the Bitcoin Improvement Proposal (BIP). A BIP serves to document modifications to the Bitcoin protocol or associated standards, providing a mechanism for peer-reviewing technical specifications, the rationale for adoption, and anticipated systemic impacts. Proponents initially present the necessity and feasibility of their ideas in open forums, such as the Bitcoin Development Mailing List (bitcoindev),[22] to gather community feedback.

Once a draft matures, it is submitted as a pull request to the GitHub BIP repository. BIP editors review the submission against formatting and scope criteria before assigning a unique number and merging it into the repository. Crucially, being registered as a BIP does not equate to formal adoption or network consensus. Final activation hinges entirely on whether developers, full-node operators, miners, wallet/exchange providers, and end-users actually implement and execute those rules in their software.

The mechanisms for altering protocol rules are broadly categorized into "softforks" and "hardforks." A softfork is a backward-compatible upgrade that introduces tighter restrictions to the existing rule set. Nodes that choose not to

upgrade can still accept blocks generated under the new rules as valid, making it easier to preserve network backward compatibility. Conversely, a hardfork introduces an upgrade that is incompatible with legacy rules. In this scenario, if all participants do not upgrade simultaneously, the network faces the permanent risk of splitting into two independent blockchains operating under divergent rules. Historically, the Bitcoin ecosystem has prioritized network stability by leaning heavily toward the gradual, conservative nature of softforks. A prime example of an upgrade successfully integrated into the mainnet through this rigorous process is Segregated Witness (SegWit), activated in 2017. Introduced via a softfork, SegWit was defined in BIP-141[23.] and separated signature (witness) data from the core transaction structure, handling it in a distinct data architecture. In essence, it reorganized the transaction blueprint by decoupling the core transactional details from the cryptographic proofs verifying them.

SegWit also changed the calculation of block capacity from a strict byte limit to a weighted metric called "Block Weight," under which witness data was discounted relative to base transaction data. This structural shift enhanced transaction flexibility and maximized block space efficiency. Ultimately, it achieved the dual benefit of maintaining flawless backward compatibility with the legacy network while effectively expanding the network's practical transaction throughput.[24.]

The activation of Taproot in 2021 marks another vital milestone in Bitcoin's evolution. Taproot introduced a highly flexible conditional spending framework and a new signature mechanism to the network.[25.] Rather than entirely deprecating the legacy ECDSA signature scheme, it introduced the option to utilize Schnorr signatures within Taproot outputs and spending paths. Schnorr signatures provide the cryptographic foundation to aggregate multi-signature transactions, making them appear as compact as single-signature inputs.

Furthermore, Taproot was designed so that even if complex spending conditions are established, only the specific path executed during the transaction is exposed to the blockchain. For instance, if multiple unlocking conditions are predefined, any unused paths remain concealed, and only the active route is written on-chain. Through this mechanism, Taproot significantly optimized data efficiency while vastly improving user privacy by preventing the exposure of unused spending conditions.

The critical takeaway here is that Taproot did not abruptly disrupt Bitcoin's existing paradigm. It preserved legacy address types and signature methods while incrementally introducing new output types and spending mechanisms. Executed as a softfork, the upgrade successfully activated at Bitcoin mainnet block height 709,632. Consequently, Taproot stands as a textbook example of how Bitcoin can safely accommodate innovative cryptographic signature schemes and complex conditional spending logic while maintaining its conservative consensus structure.

The achievements of SegWit and Taproot demonstrate that Bitcoin is not a static, unchanging system, but a living system capable of deliberate, organic evolution. SegWit transformed the paradigm of transaction structures and block space utilization, while Taproot introduced next-generation signature schemes and script expressions. Both major upgrades were achieved through prolonged verification and consensus within Bitcoin's inherently conservative

development culture. This historical precedent provides invaluable context for understanding the forthcoming post-quantum migration. Bitcoin does not change haphazardly or impulsively, but it possesses the resilience to evolve through consensus when threats and necessities are clearly proven.

Technical Paths for a Gradual Transition

The initial phase of a quantum-resistant migration does not involve the overnight deprecation of the legacy signature scheme. Rather, it requires establishing a technical gateway capable of accommodating new spending conditions within the Bitcoin protocol. As demonstrated by earlier upgrades such as SegWit and Taproot, the transition to post-quantum cryptography will likely begin with the introduction of new output structures, allowing users to voluntarily migrate their assets to secure addresses over an extended period rather than immediately invalidating existing addresses and transaction models.

One of the leading proposals currently under discussion is outlined in BIP-360, which introduces a new output type known as Pay-to-Merkle-Root (P2MR).[26.] Taproot outputs do not record an internal public key directly. Instead, they record an output key derived from the internal public key and script-tree information. P2MR, by contrast, removes Taproot's key-path spending mechanism and records only the Merkle root of a script tree containing multiple spending conditions. Put simply, rather than revealing in advance the public keys or detailed conditions required to spend an asset, P2MR leaves only a cryptographic commitment on the blockchain indicating that those conditions exist.

Because P2MR records the script tree as a 256-bit hash commitment, it can provide an output structure capable of accommodating future quantum-resistant signature schemes and related script conditions should they eventually be introduced into Bitcoin.

BIP-360 distinguishes between two categories of quantum attacks: long-exposure attacks and short-exposure attacks. A long-exposure attack refers to a scenario in which a public key remains visible on the blockchain for an extended period, allowing an attacker to use a future quantum computer in an attempt to recover the corresponding private key. Structures such as Pay-to-Public-Key (P2PK) and Taproot, both of which depend on public-key-based spending mechanisms, may be relatively more vulnerable to this type of attack.

A short-exposure attack, by contrast, involves deriving a private key from a public key that has been revealed while a transaction is sitting in the mempool and then broadcasting a competing transaction with a higher fee. P2MR still requires the public key associated with the active signature scheme to be revealed when funds are actually spent. Consequently, in systems such as ECDSA or Schnorr signatures, where public keys become visible at the time of spending, P2MR alone cannot prevent short-exposure attacks that occur while a transaction remains in the mempool.

BIP-360 is not, in itself, a proposal to introduce quantum-resistant signature algorithms such as ML-DSA or SLH-DSA into Bitcoin. Rather, it is best understood as a proposal to establish an effective output structure that would

enable the future adoption of quantum-resistant signature schemes. The BIP-360 document itself notes that P2MR alone cannot prevent short-exposure attacks targeting public keys that are briefly exposed in the mempool before transaction confirmation. It further acknowledges that more sophisticated quantum threats may ultimately require the adoption of quantum-resistant signature schemes.

In this sense, P2MR is less a complete defense than a foundation for transition. It preserves as much as possible of the flexible script structure provided by Taproot while reducing the problem of long-term public-key exposure before coins are spent. It also represents a conservative approach aimed at minimizing the changes required of existing infrastructure—including wallets, exchanges, and software libraries—by leveraging the current Taproot framework wherever possible.

This approach, however, is not without costs. Because P2MR eliminates Taproot's key-path spending mechanism, transaction data can become larger than in the simplest Taproot key-path spends. A standard Taproot key-path spend requires only a signature, whereas P2MR also requires the spending script and proof that the script is included in the Merkle tree.[27]

Consequently, P2MR reduces the risk associated with the long-term exposure of public keys, but it may impose additional costs in terms of block space and transaction fees. This trade-off becomes even more significant if quantum-resistant signature algorithms are eventually introduced into Bitcoin. In general, quantum-resistant signatures are substantially larger than ECDSA or Schnorr signatures, meaning that improved security may come at the cost of increased transaction size and verification overhead.

Nevertheless, the broader significance of BIP-360 lies in demonstrating that Bitcoin's transition to a quantum-resistant system does not necessarily require a disruptive hard fork. By preserving existing address formats while introducing new output types, P2MR provides a voluntary migration path through which users can gradually move their assets to more secure addresses according to their own risk assessments. In this framework, the transition is not defined as a single mandatory event but as a gradual process in which secure pathways are established first and migration follows over time.

In summary, Bitcoin's transition to a quantum-resistant system is not a technological impossibility. The broader cybersecurity ecosystem has already entered the standardization and deployment phase of post-quantum cryptography. Within Bitcoin, active efforts are underway to develop structures that minimize public-key exposure and accommodate future quantum-resistant signature schemes.

Moreover, a post-quantum transition differs fundamentally from governance disputes such as changing Bitcoin's fixed supply cap, where stakeholder interests may be directly opposed. If quantum computing advances to the point where Bitcoin's signature schemes become vulnerable, miners, developers, exchanges, wallet providers, and investors alike would face a common systemic threat. Under such circumstances, broad consensus regarding the necessity of a transition would likely emerge relatively quickly.

The central challenge therefore lies not in whether a transition is possible, but in how it is implemented. Success will depend on a complex combination of

factors: selecting appropriate quantum-resistant signature algorithms and output structures, deploying them through a stable soft fork, securing support across wallets, exchanges, and infrastructure providers, and encouraging users to migrate their assets.

Technical pathways can be designed, but they only become effective through voluntary adoption and decentralized coordination across the network. At this point, the discussion naturally turns to two additional challenges: the timeline of the transition itself and the question of how to deal with coins that do not—or cannot—participate in that transition. Depending on when a quantum threat emerges and how quickly the ecosystem responds, the process may unfold as an orderly technological evolution or as a race against time. The treatment of these non-migrating coins therefore becomes one of the most important unresolved questions in Bitcoin's post-quantum future.

The Transition Timeline and the Dilemma of Residual Coins

To comprehensively evaluate Bitcoin's post-quantum migration process, one must first critically examine the dimension of the "timeline." The threat levied against the Bitcoin network by quantum computing will not materialize as a single, uniform event. More precisely, the nature of the threat shifts dramatically depending on when the protocol transition is initiated. A scenario where Bitcoin proactively integrates defensive mechanisms prior to the commercialization of quantum attacks is fundamentally distinct from a scenario where remediation begins reactively after quantum exploits have surfaced. The former constitutes a manageable, orderly transition, whereas the latter escalates into a high-stakes race against adversarial exploitation.

Under the proactive scenario, the transition can be engineered as a controlled, non-disruptive process that preserves network stability. The network can integrate post-quantum address formats and advanced spending conditions into the protocol well in advance, allowing users an extensive grace period to securely relocate their assets. Here, the primary objective is ensuring procedural simplicity and accessibility, enabling holders to utilize their private keys to migrate assets to post-quantum addresses. Users execute standard transactions via upgraded wallet or exchange interfaces, and network nodes validate these inputs against the updated rule set. Within this paradigm, the transition ceases to be a systemic shock and instead operates as a prolonged, predictable asset migration.

Conversely, a reactive scenario introduces extreme security risks. If protocol upgrades are only initiated after quantum computers reach the technical threshold required to derive private keys from on-chain public keys, the very act of moving legacy assets to secure post-quantum addresses becomes a vector for exploitation. In a normal environment, a user simply broadcasts a transaction to route assets to a new address. However, in an ecosystem haunted by an active quantum adversary, the moment a user signs a transaction to spend from a legacy output, their public key is exposed to the mempool. A sophisticated attacker can intercept this brief exposure, compute the private key almost instantly, and broadcast a front-running transaction with

a higher fee to seize the assets before the original transaction is confirmed.

At this stage, the transition ceases to be an orderly software deployment; it becomes a race over whose transaction gets confirmed first. Ultimately, the timing of the transition is the definitive variable that dictates whether this challenge culminates in an orderly evolution or a destructive conflict.

This precise problem statement is consistently mirrored in academic research focused on post-quantum migration protocols. In a 2018 paper, a research team led by I. Stewart at Imperial College London's Centre for Cryptocurrency Research and Engineering proposed a transition protocol utilizing a "Commit-Delay-Reveal" mechanism.[28] This architecture was specifically designed for an adversarial environment where post-quantum signature schemes have been integrated into the Bitcoin protocol, yet a substantial volume of assets remains trapped in legacy, non-quantum-resistant outputs. True to its name, the protocol enforces three distinct phases:

- **Commit**: In the initial phase, a user does not expose their legacy public key or their new post-quantum public key directly. Instead, they publish a single cryptographic hash that binds both keys together. Because hash functions are resistant to reverse-engineering, an adversary cannot work backward from the hash to discover the public keys. This hash serves as an on-chain commitment—a sealed promise of a future claim without revealing the underlying data.
- **Delay**: Following the hash commitment, the user cannot immediately move the assets. Instead, they must wait for a required period (the researchers proposed an illustrative period of six months). This mandatory delay ensures that the commitment becomes deeply embedded in the blockchain's history. If the commit were only a few blocks deep, a well-resourced attacker could execute a chain reorganization to erase the record and substitute a fraudulent one. However, once the commitment is buried under months of accumulated blocks, the economic and computational cost of rewriting the ledger becomes prohibitively high, making the commitment practically irreversible.
- **Reveal:** After a sufficient delay period has elapsed, the user reveals the legacy public key, the new quantum-resistant public key, and the verification data proving that the relationship between the two keys was previously committed. At the same time, the user submits a signature generated with the new quantum-resistant private key to demonstrate control of the new quantum-resistant key as well. Even if an attacker observes the revealed legacy public key and subsequently derives the corresponding legacy private key, they still do not possess the new quantum-resistant private key. Nor can they retroactively create the commit record that only the legitimate owner could have established before the delay period began. As a result, under the new rules, learning the legacy private key after the fact is not sufficient to seize the asset. The authors argue that this approach can remain effective even in a scenario where ECDSA has already been compromised. Although it requires modifications to the Bitcoin protocol, they note that it could be implemented through a soft fork.

Of course, this architecture is not a panacea for all security blind spots. The

mechanism primarily safeguards outputs whose public keys have not yet been directly exposed on the blockchain—legacy P2PKH addresses being a prime example. It cannot fully remediate legacy outputs whose public keys are already deeply woven into the blockchain's historical ledger, nor can it accommodate crisis scenarios requiring the rapid mass evacuation of assets under time constraints.

Nevertheless, the profound takeaway from this research is clear: the core of the post-quantum transition is not merely the introduction of a new cryptographic signature scheme. The authentic challenge lies in engineering the institutional and technical procedures required to safely usher legacy assets into a new security paradigm. Deciding when public keys are revealed, determining the duration of commit delays, and establishing the validation criteria for legitimate transfers all require a fundamental rewriting of consensus rules. In short, post-quantum migration is not an isolated swap of cryptographic algorithms; it is the comprehensive design of a governance architecture to relocate legacy capital into a new security order.

Yet, regardless of how flawlessly a migration architecture is designed, not all on-chain assets will participate in this migration. Legacy software wallets may lack the updates required to compile new address formats, users may remain oblivious to the urgency of the migration, and long-term passive investors may consciously choose inaction.

Far more severe categories of assets exist: coins where private keys have been permanently lost, owners who have passed away, or accounts where physical access has been completely severed. These "lost coins" cannot move autonomously, irrespective of how innovative the transition pathways are. Ultimately, technical migration vectors can only rescue assets backed by active human agency; they cannot inherently dissolve the systemic risks posed by stagnant, unmanaged assets anchored to the ledger.

Furthermore, relying on a hardfork-centric overhaul does not bypass this dilemma. Proponents of a hardfork possess the authority to fundamentally rewrite the treatment of the existing UTXO set. They can theoretically design an architecture that freezes high-risk addresses, imposes punitive conditions on legacy withdrawals, or alters asset weights during the generation of a new chain snapshot.

However, implementing heavy-handed consensus rules from the top down does not automatically relocate unmanaged legacy assets into post-quantum secure addresses. Attempting to forcefully integrate legacy coins into a post-quantum architecture without an explicit signature from the rightful owner requires the enactment of artificial intervention rules that violate the foundational principles of a decentralized ledger. Thus, a hardfork is not a magical solution that dissolves the dilemma; rather, it is a volatile governance arena that demands the ecosystem make disruptive political judgments regarding which assets to legitimize and which to abandon.

At this point, the transition to quantum-resistant cryptography effectively divides bitcoin into two categories: coins that can be moved and coins that remain unmoved or cannot be moved. The former category is relatively straightforward. If owners retain their private keys, wallets and exchanges support the new address format, and sufficient time is available, these coins

can be migrated to the new system. The challenge lies with the latter category. Legacy coins whose public keys have already been exposed, coins that have remained dormant for extended periods, lost coins whose private keys are no longer available, and assets whose owners' intentions cannot be verified—such as Satoshi Nakamoto's early coins—remain separate issues that must be addressed during the transition process.

Consequently, in Bitcoin's transition to a quantum-resistant system, developing new signature schemes and address formats is only the first task. The more difficult question is how to regard coins that do not migrate to the new system. Are they the assets of legitimate long-term holders? Are they inaccessible lost coins? Are they vulnerable assets that could be targeted by quantum attackers because their public keys have been exposed? Or are they risk-bearing assets that the network may need to address in a different way? This is precisely the issue explored in the next chapter. The question is not whether Bitcoin can transition to a quantum-resistant framework, but how to deal with coins that do not—or cannot—participate in that transition.

05

The Problem of Satoshi Coins and Lost Coins

As demonstrated in the preceding chapter, the most formidable challenge of a post-quantum transition does not conclude with the mere provisioning of a next-generation address framework. This chapter examines the profound dilemmas caused by assets that either fail or are unable to participate in this migration—specifically, legacy coins with historically exposed public keys, the "Satoshi coins," and permanently lost coins.

Coins with Historically Exposed Public Keys

From the perspective of quantum vulnerability, not all Bitcoin Unspent Transaction Outputs (UTXOs) share an identical risk profile. The critical determinant resides in whether a public key is already visibly exposed on the distributed ledger, and the duration for which that exposure has persisted. In the nascent stages of Bitcoin, the Pay-to-Public-Key (P2PK) script format was predominantly utilized. Under P2PK, the public key is directly embedded within the locking script (scriptPubKey). Consequently, from the exact moment of output creation, the public key is exposed on the blockchain. Should future quantum computers achieve sufficient logical qubit capacity, these specific UTXOs will become the primary targets for long-exposure attacks, as adversaries can attempt to derive the corresponding private keys directly from these visible public keys.

Conversely, under the Pay-to-Public-Key-Hash (P2PKH) framework, which has been widely adopted in subsequent years, it is not the raw public key but its cryptographic hash that is recorded on-chain. Within this architecture, the public key remains completely concealed until the asset is actively spent. Because a public key hash represents an additional layer of cryptographic compression, it is computationally infeasible for an attacker to deduce the public key from the hash alone. In this regard, P2PKH exhibits a significantly lower vulnerability to long-exposure quantum attacks compared to its predecessor, P2PK.

Nevertheless, P2PKH does not afford absolute immunity against quantum threats. The vulnerability manifests the precise instant a UTXO is spent, as the transaction must expose the raw public key to the network for verification. Under ordinary circumstances, this brief exposure carries nominal risk; once the transaction is validated and finalized, the target UTXO is entirely consumed. If no residual balance remains at that address, an adversary deriving the private key gains nothing. However, the paradigm shifts

dramatically when addresses are reused. If a public key is revealed via an initial expenditure, and residual funds remain at that same address—or if the address subsequently receives new inflows—an adversary gains a prolonged window of opportunity to execute private key derivation attacks against the now-exposed public key.

For this reason, avoiding address reuse has long been championed as a foundational security best practice within the Bitcoin ecosystem. In the shadow of emerging quantum threats, this principle transitions from a standard privacy recommendation to a critical defensive strategy designed to compress the exposure window of public keys. Ultimately, coins with unexposed public keys and those with historically visible public keys occupy fundamentally disparate risk domains, and any rigorous blueprint for a post-quantum transition must proceed from this baseline divergence.

[Table 5] **Quantum Vulnerability by Bitcoin Address Type**

Address / Output Type	How the Public Key is Exposed	Quantum Risk
P2PK	Public key is directly exposed when the output is created.	Highest risk of long-exposure attacks.
P2PKH	Only the public key hash is visible prior to spending; the public key is revealed when spent.	Risk increases significantly with address reuse.
P2TR	The Taproot output key is exposed on-chain.	May be vulnerable to long-exposure attacks.
P2MR	Only the Merkle Root is recorded in the output.	Effectively mitigates long-term public key exposure; however, short-exposure risks at the time of spending require separate countermeasures.

Satoshi Coins and Lost Coins

This structural bifurcation becomes particularly relevant when analyzing Bitcoin's oldest, dormant allocations. During the network's infancy, P2PK outputs—which directly expose the public key—were standard, and a significant portion of these funds has remained entirely motionless to this day. Most notable among these are the approximately 1.0 to 1.1 million bitcoins[29] estimated to have been mined by Satoshi Nakamoto, an allocation that serves as a recurring focal point in post-quantum security discourse.

The gravity of the Satoshi coins dilemma is self-evident. Should a large amount of bitcoin associated with the creator suddenly move via quantum exploitation, the global market would unlikely perceive it as a localized security breach. Instead, it would be interpreted as a signal that Bitcoin's oldest cryptographic locks have been compromised, compounded by the sudden materialization of a large potential supply overhang. Representing roughly 5% of the total 21-million hard-capped supply, the mere perception that this supply has become liquid or vulnerable to market dumping could deliver a severe shock to Bitcoin's valuation and systemic credibility.

Yet the dilemma extends beyond the creator's holdings. Permanently inaccessible assets—commonly known as "lost coins," where private keys have been lost due to negligence, hardware failure, or the owner's death—present a separate challenge in a post-quantum migration. On-chain inactivity alone cannot tell us whether a coin is truly lost or simply being held long-term. The real problem lies in coins whose private keys are genuinely lost or whose owners can no longer demonstrate their intent. Because these coins cannot migrate to a new post-quantum address, they leave an unresolved question about how the network should handle them.

Consequently, the post-quantum transition simultaneously exposes two distinct vectors of systemic friction. The first is a cryptographic vulnerability concerning historically exposed public keys, which become immediate targets for exploitation once sufficiently powerful quantum computing architectures emerge. The second is a governance dilemma regarding the treatment of stagnant or unmanageable assets. If a sovereign holder maintains control over their private keys, migration to a secure domain is a matter of voluntary execution. However, assets stripped of human agency or missing their underlying private keys are structurally precluded from participating in the upgrade. While these two problems frequently intersect, their underlying characteristics are profoundly different: the former is a technical problem of cryptographic vulnerability, whereas the latter is a political and philosophical judgment regarding ownership, immutability, and the legitimate boundaries of network intervention.

This combination of vulnerabilities forces the Bitcoin network to interpret the true nature of motionless capital, transitioning the debate into a conflict between "quantum theft" and "quantum recovery." If an adversarial actor utilizes a quantum computer to expropriate an old coin with an exposed public key, should the decentralized ledger accept the transaction as a valid expenditure under its procedural rules? Alternatively, to preempt such exploitation, should the network enforce strict protocol-level restrictions on legacy spending after a designated grace period has lapsed? This inquiry transcends simple security policy; it strikes at the core of how ownership is defined within Bitcoin, and the extent to which the network can legitimately manipulate established consensus rules in the name of security.

Four Paradigms Surrounding Stagnant Capital

The governance choices available to the ecosystem can be categorized into four primary paradigms.

The first paradigm is the Status Quo. Under this approach, ownership within Bitcoin is defined strictly by the ability to produce a valid signature, regardless of identity, intent, or how long the coins have been held. Even if a public key is exposed on-chain and someone uses a quantum computer to compute the corresponding private key and produce a valid signature, the network does not block the transaction. Under this approach, a quantum-driven transaction is not seen as an external attack, but rather as a natural consequence of the cryptographic foundation becoming no longer sufficiently difficult to break. The advantage is clear: the network does not judge ownership after the fact, does

not arbitrarily block specific coins, and maintains the rule that 'a valid signature is a valid spend.' However, the cost is significant. From the original owner's perspective, the transaction amounts to theft. Furthermore, if Satoshi coins or large long-dormant holdings suddenly move, the market may interpret this as a signal that a massive potential supply has been unlocked. The rules are preserved, but the economic trust those rules were meant to protect may be shaken.

The second paradigm is Freezing. This approach restricts the spendability of outputs with exposed public keys that fail to migrate to post-quantum addresses after a designated grace period, effectively making certain legacy UTXOs unspendable. Freezing is proposed as a safeguard to prevent a quantum attacker from computing private keys and moving large amounts of assets. Given that the sudden movement of Satoshi's coins or long-dormant outputs could trigger severe market panic even without actual selling, proactive freezing appears to be the most direct defensive measure. However, freezing directly conflicts with Bitcoin's foundational principles. In Bitcoin, control over assets is defined by possession of a private key, not by legal identity or the approval of a central authority. If the network blocks a coin from being spent simply because it has been "dormant for a long time" or is "vulnerable to quantum attack," this creates an exception to the principle that anyone who can produce a valid signature can spend the coin. Furthermore, legitimate long-term holders, individuals who recovered their keys late, or users who missed the migration window due to technical difficulties could also find their access blocked. Freezing may deter attackers, but it carries the risk of locking out rightful owners at the same time.

The third paradigm is Conditional Recovery, a middle ground between preserving existing rules and full freezing. Under this approach, once the migration grace period expires, coins with exposed public keys can no longer be spent using standard elliptic curve signatures alone. Instead, legitimate owners are given an alternative pathway to prove they have controlled the asset and transfer it to a post-quantum address. Theoretically, Zero-Knowledge Proofs (ZKPs) could be used to provide a limited recovery path for certain types of owners.[30] A Zero-Knowledge Proof is a cryptographic technique that allows someone to prove they know a secret without revealing the secret itself. It may therefore be possible to design a recovery process where an owner proves legitimate control under certain conditions without exposing any private information on the blockchain.

However, this approach must be handled with care. If the recovery process relies solely on proving knowledge of the private key, it cannot distinguish between the rightful owner and a quantum attacker who has computed the same key. For conditional recovery to be meaningful, it requires a separate secret that has not been exposed on the blockchain—something that cannot be reverse-engineered from the public key and is known only to the legitimate owner. A BIP-39[31] seed phrase, used by many modern wallets—a sequence of words used to recover a wallet—is one such example. A quantum attacker cannot derive a seed phrase from an exposed public key. However, early Bitcoin wallets did not use seed phrases, and there is no way to verify whether any such separate secret exists for Satoshi coins or other early mining outputs. This approach also offers no solution for genuinely lost coins where the private key no longer exists. Therefore, while ZKP-based conditional recovery may

offer a viable path for some modern wallets and users, it cannot serve as a universal solution to the broader problem of Satoshi coins and lost coins.

The fourth paradigm is Burning, a more extreme choice than freezing. Under this approach, any output with an exposed public key that fails to migrate to a post-quantum address within the grace period is permanently made unspendable or removed from the circulating supply. Burning eliminates the risk of a quantum attacker seizing lost or dormant coins and introducing them to the market. In the long run, it also removes the uncertainty that quantum-vulnerable coins could appear on the market at any time. However, the cost is the greatest of all four options. Burning sets a precedent that the network can arbitrarily invalidate the legitimate long-term holdings of rightful owners. If one of the reasons Bitcoin is trusted is its resistance to arbitrary confiscation and censorship, burning is the option that conflicts most directly with that principle. The moment a precedent is set that part of the supply can be removed for security reasons, Bitcoin's procedural immutability can no longer carry the same meaning it once did.

Ultimately, each of the four options carries its own risks. The status quo preserves procedural neutrality but accepts the risk that a quantum attacker could move old coins. Freezing blocks that possibility but may also restrict the spending rights of legitimate long-term holders. Conditional recovery attempts a compromise but only works in limited cases where a separate proof condition can distinguish the rightful owner from a quantum attacker. Burning is the strongest defensive measure but does the most damage to Bitcoin's censorship resistance and the principle that ownership is inviolable.

Therefore, the problem of Satoshi coins and lost coins is not a question of which cryptographic algorithm to adopt when quantum computers arrive. It is a question of how the Bitcoin ecosystem chooses to interpret coins that do not move. Are these the legitimate assets of long-term holders? Are they inaccessible lost coins? Are they ownerless assets that a quantum attacker could seize? Or are they risky assets that the network must protect or isolate? These questions cannot be answered by technology alone. They require the Bitcoin community to make a judgment about what it values most—ownership, immutability, security, or economic stability.

This is where the next chapter begins. If no consensus is reached on these options, Bitcoin may face not just a technical upgrade but a deeper conflict over the legitimacy of the chain and the meaning of its procedural immutability.

06

Hardforks and the Forking of Legitimacy

Two Bitcoins That a Hardfork Could Produce

If the four choices discussed in the previous chapter fail to reach a single consensus, the issue ultimately shifts to determining which rules to follow. In this situation, a hardfork is not merely an upgrade method, but an event where a network that shares the same history splits into chains operating under different rules. Looking at the exact same UTXO set, one chain can continue to recognize the existing signature rules as they are, while the other can place restrictions on spending coins that are vulnerable to quantum attacks. In this scenario, Bitcoin faces a choice between two chains with different interpretations of ownership, rather than a simple technical upgrade.

One side may choose to maintain the existing rules. On this chain, as long as a valid signature is presented, no distinction is made between whether the signature was created by the original owner or with a private key computed by a quantum computer. It maintains the principle that ownership in Bitcoin is determined by a valid signature, not by identity or intent. From this perspective, the network's role is not to interpret who the owner is, but to verify signatures that match the rules.

Conversely, the other chain could place separate restrictions on old coins with exposed public keys or coins that have not moved for a long time. It could restrict the spending of coins that do not migrate to post-quantum addresses within a certain grace period, require additional conditions for coins with exposed public keys, or freeze or burn specific categories of UTXOs. This chain judges that the existing signature rules alone are no longer sufficient. In a situation where quantum computers undermine the cryptographic foundations, a valid signature alone cannot determine rightful ownership.

At this point, the two chains share the same history but implement different philosophies. The former maintains the principle that "a valid signature constitutes a valid spend," while the latter holds that "under conditions where quantum attacks are possible, ownership cannot be judged by the existing signature rules alone." The former emphasizes procedural neutrality, while the latter emphasizes system protection and economic stability. It is difficult to say that either side is simply right or wrong, because both hold onto core values of Bitcoin.

A similar case already exists in cryptocurrency history. In 2016, within the Ethereum ecosystem, a smart contract project designed like a decentralized investment fund, "The DAO," was attacked. The attacker exploited a

vulnerability in The DAO's smart contract, rather than the Ethereum protocol itself, to extract a large amount of Ether. This raised the question of whether outcomes produced by code should be accepted as they are, or whether the community should intervene to reverse the damage. Ultimately, the Ethereum community chose a chain that recovered the funds through a hardfork.[32] This hardfork was executed at block 1,920,000, and the Ethereum Foundation explained at the time that approximately 12 million ETH were moved to a withdrawal contract. Some participants who opposed this maintained the existing chain, resulting in the formation of "Ethereum Classic (ETC)." This event shows that a hardfork is not just a technical fix, but a question of what should be viewed as legitimate history.

The quantum computing issue raises the same question: which transactions should be recognized as legitimate history. If a consensus cannot be reached on whether the movement of old coins with exposed public keys should be viewed as a valid spend or restricted as a quantum theft, Bitcoin will split into chains with different ownership philosophies. In this case, which chain is the "legitimate Bitcoin" is not decided by code alone. Legitimacy is formed through the choices of node operators, miners, exchanges, wallet providers, custody institutions, institutional investors, self-custody users, and ordinary users—choices about which chain to mine, which to label as BTC, which to recognize as a custody asset, and which to hold long-term. A hardfork is an event that splits code, but legitimacy is created through the social choices that run that code and give it value.

Institutional Hardforks and the Allure of Public Goods

A more radical possibility can also be considered here. Instead of freezing or burning old coins within the existing Bitcoin chain, this involves creating a chain with entirely new rules and encouraging institutional infrastructure to recognize that chain as the new standard Bitcoin chain. In this scenario, a snapshot would be taken based on a specific block height, and coins that do not move to post-quantum addresses within a certain grace period would be classified as "unmoved assets." The new chain could then be designed to allocate these coins to a network security fund, a development fund, or a public goods fund.

Technically, such a fork is not impossible. Since a hardfork introduces new consensus rules that are incompatible with existing rules, the new chain can be designed to handle UTXOs that do not meet certain conditions differently from the existing chain. For example, a rule could be created where coins with exposed public keys that do not move to post-quantum addresses by a certain point are no longer recognized as assets that individuals can spend, and are instead assigned to the new chain's public goods fund.

In this case, supporters of the fork could argue that rather than leaving old, vulnerable coins as easy targets for quantum attackers, it is more rational to quarantine them on the new chain and convert them into resources for the security and development of the entire network. As discussed in the previous chapter, the possibility of Satoshi coins or long-dormant coins moving can be perceived by the market as a potential supply shock. In this respect, a

redistributive hardfork is not just a technical defense, but a political and economic proposal to redefine the ownership and use of old coins in light of the quantum threat.

However, the core of this scenario is legitimacy, not technology. Anyone can create a forked chain with new rules, but no fork automatically gains Bitcoin's legitimacy. For the new chain to acquire the name and liquidity of Bitcoin, major exchanges, custody providers, ETF issuers, institutional investors, miners, node operators, and self-custody users must accept that chain. Especially since Bitcoin has integrated deeply into institutional finance, which chain is recognized as the underlying asset for ETFs, which chain is treated as an eligible asset in custody systems, and which chain maintains the BTC ticker on major exchanges can heavily influence market choices.

Institutional products like spot Bitcoin ETFs may include provisions that leave the judgment of which network to view as appropriate for trust purposes to the sponsor's discretion in the event of a hardfork. For example, the prospectus for BlackRock's iShares Bitcoin Trust ETF explains that if a hardfork occurs on the Bitcoin network, the sponsor can use its discretion, within the scope permitted by the trust agreement, to determine which network is appropriate for trust purposes.[33.]

This does not mean that the ETF sponsor can change the consensus rules of the Bitcoin network. It simply means that if multiple chains coexist due to a hardfork, the sponsor can decide which chain to treat as Bitcoin within the financial product it manages. In other words, institutional products can become selectors that grant a specific chain its name and liquidity in a fork scenario, rather than mere observers.

This is precisely where the political and economic significance of an institutional hard fork emerges. If large financial firms such as BlackRock, major exchanges, custodians, and even regulators and governments effectively favor a particular forked chain through tax laws, accounting standards, custody rules, and exchange regulations, many institutional investors and financial products may have strong incentives to follow that chain. In that case, a hardfork is no longer merely a technical split within the community. Instead, it becomes a question of which chain institutional infrastructure chooses to recognize and support with liquidity. Bitcoin's legitimacy will be contested not only in code repositories and developer mailing lists, but also through ETFs, custodial services, accounting standards, tax rules, exchange listings, and institutional investment portfolios.

Yet at that very point, a redistributive hardfork most directly challenges Bitcoin's identity. Just because institutional entities support a specific fork does not automatically mean there is consensus across the entire Bitcoin community. If self-custody users, full-node operators, miners, and the existing developer community do not accept those rules, the new chain may remain a separate forked chain approved by institutions rather than a continuation of Bitcoin itself. Institutions can move a significant portion of the name and liquidity, but that alone cannot completely determine Bitcoin's legitimacy.

A more fundamental issue is that the justification of public goods can conflict with Bitcoin's ownership principles. The idea of securing unmoved coins as public goods runs into the problem that lost coins and long-term held coins are

indistinguishable on-chain. Just because a coin has not moved for years or decades does not mean it is definitively a lost coin. It could be the asset of a long-term holder, an asset waiting to be inherited, or an asset unmoved for political, legal, or personal reasons. Therefore, incorporating a coin into a public goods fund simply because it has not moved by a certain point requires accepting the burden of re-judging ownership, going beyond a simple security measure.

Who will manage the public goods fund is also a question. If it is a developer fund, which developers will distribute those resources? If it is a security fund, which miners or validation infrastructure will receive the rewards? If a government manages it, Bitcoin risks being drawn back into the very sovereign financial system it was designed to operate outside of. If private financial institutions manage it, is it truly a public good, or is it a new reserve controlled by institutional custody infrastructure? The moment these questions are opened, Bitcoin shifts back to a system that requires trusted managers.

Therefore, a redistributive hardfork is both a powerful scenario for institutionally managing the risks of quantum attacks and the option that most directly tests Bitcoin's censorship resistance, the inviolability of ownership, supply rules, and the source of its legitimacy. It appears to be a proposal to convert lost coins into public goods, but in reality it raises deeper questions: who can judge what constitutes a lost coin, who can decide what Bitcoin is, and who holds the authority to redistribute unmoved assets. These questions cannot be answered by technology alone. Ultimately they come down to a single question: will Bitcoin's future be shaped by institutions and governments, or by the decentralized community of users, developers, and node operators who have always been its true foundation?

Conclusion

The quantum computing issue is a question that exposes Bitcoin's vulnerabilities while at the same time reaffirming what it stands for. Bitcoin's total supply is hard-capped at 21 million, its halving schedule is predictable, and there is no institution—like a central bank or government—that can adjust monetary rules at its discretion. This predictability is a key reason Bitcoin has built long-term trust. If the rules of a currency change easily, it is difficult for anyone to trust it as a long-term store of value.

However, Bitcoin's immutability does not mean that every part of its code is frozen forever. Bitcoin has already been upgraded several times and may need to change in the future to respond to new threats. What matters is not whether it changes, but how it changes. Parts of the code can change in response to new threats, but for that change to gain legitimacy, it must pass through public verification, voluntary adoption, and distributed consensus. Bitcoin's immutability does not mean that nothing changes; it means that no one can change the rules without consensus.

If a situation arises where the cryptographic system is genuinely compromised, maintaining the existing rules could instead put the entire system at risk. If signature rules designed to protect ownership no longer keep ownership secure, adjusting those rules is not an abandonment of principle—it is a change made in service of it. In this respect, the post-quantum transition does not conflict with Bitcoin's immutability; it is an evolution necessary to sustain it.

However, not all changes carry the same legitimacy. The introduction of new signature schemes and address formats can be accepted as an upgrade that strengthens security. But the handling of unmoved coins, hardforks, institutional chain selection, and attempts to convert assets into public goods raise deeper questions. It is not simply a matter of which technology to adopt, but a question of who can judge ownership, who can decide what Bitcoin is, and what kind of process makes a change legitimate.

Therefore, the fundamental question raised by quantum computing is whether Bitcoin can handle necessary changes without undermining its own principles. When change is necessary, it must be possible. But no one should be able to change it arbitrarily. That tension—between the ability to adapt and the resistance to arbitrary change—is where Bitcoin's future will be written.

Notes

1. Matt Clinch, "Bitcoin Hacked: Price Stumbles After Buying Frenzy," CNBC, April 4, 2013, accessed February 28, 2026, https://www.cnbc.com/2013/04/04/bitcoin-hacked-price-stumbles-after-buying-frenzy.html

2. Yoshifumi Takemoto and Sophie Knight, "Mt. Gox Files for Bankruptcy, Blames Hackers for Losses," Reuters, February 28, 2014, accessed March 13, 2026, https://www.voanews.com/a/reu-mt-gox-files-for-bankruptcy-blames-hackers-for-losses/1861341.html

3. Stan Higgins, "The Bitfinex Bitcoin Hack: What We Know (and Don't Know)," CoinDesk, August 4, 2016, accessed February 28, 2026, https://www.coindesk.com/markets/2016/08/03/the-bitfinex-bitcoin-hack-what-we-know-and-dont-know

4. No confirmed case exists in which Bitcoin's core cryptographic mechanisms or consensus system has been compromised by an external attack, resulting in the theft of assets at the network level. However, critical vulnerabilities have been discovered in the code. In August 2010, a transaction exploiting an integer overflow bug was included in block 74,638, creating approximately 184.4 billion bitcoins. Developers quickly released a patched version, and the network abandoned the affected branch and reorganized onto the valid chain. In 2018, CVE-2018-17144 was disclosed. The vulnerability enabled denial-of-service (DoS) attacks through duplicate-input handling and, more seriously, could have allowed inflation of the bitcoin supply. Fortunately, a patched release was issued before the vulnerability could be exploited. Together, these incidents demonstrate not only that Bitcoin is not a flawless system, but also how an open-source consensus protocol can identify vulnerabilities, address them, and continuously improve itself.
"Bitcoin Block 74,638," Blockchain.com Explorer, mined August 15, 2010, https://www.blockchain.com/explorer/blocks/btc/74638.
Bitcoin Core, "Disclosure of CVE-2018-17144," September 20, 2018, https://bitcoincore.org/en/2018/09/20/notice.

5. Strictly speaking, valid Bitcoin private keys consist of numbers within the range permitted by secp256k1, making their total number slightly smaller than 2^{256}. However, the difference is negligible relative to the overall scale, so approximating the key space as 2^{256} is generally acceptable.

6. Following the Taproot upgrade in 2021, Bitcoin gained the ability to use Schnorr signatures for Taproot outputs. Like ECDSA, Schnorr signatures are ECC-based signature schemes that operate on the secp256k1 elliptic curve and therefore share the same class of vulnerabilities with respect to quantum attacks.

7. Peter W. Shor, "Algorithms for Quantum Computation: Discrete Logarithms and Factoring," in Proceedings of the 35th Annual Symposium on Foundations of Computer Science (FOCS, 1994), 124-34, https://doi.org/10.1109/SFCS.1994.365700.

8. Shor, "Algorithms for Quantum Computation," 124-34.

9. Lov K. Grover, "A Fast Quantum Mechanical Algorithm for Database Search," in Proceedings of the 28th Annual ACM Symposium on Theory of Computing (1996), 212-19, https://doi.org/10.1145/237814.237866.

10. Jay Gambetta, "The Hardware and Software for the Era of Quantum Utility Is Here," IBM Quantum Research Blog, December 4, 2023, accessed May 21, 2026, https://www.ibm.com/quantum/blog/quantum-roadmap-2033.

11. Ryan Mandelbaum et al., "How IBM Will Build the World's First Large-Scale, Fault-Tolerant Quantum Computer," IBM Quantum Research Blog, June 10, 2025, accessed May 21, 2026, https://www.ibm.com/quantum/blog/large-scale-ftqc.

12. Hartmut Neven, "Meet Willow, Our State-of-the-Art Quantum Chip," Google Blog, December 9, 2024, accessed February 28, 2026, https://blog.google/technology/research/google-willow-quantum-chip.

13. Catherine Bolgar, "Microsoft's Majorana 1 Chip Carves New Path for Quantum Computing," Microsoft Source, February 19, 2025, accessed May 21, 2026, https://news.microsoft.com/source/features/innovation/microsofts-majorana-1-chip-carves-new-path-for-quantum-computing.

14. Kevin Williams, "What Google's Quantum Computing Breakthrough Willow Means for the Future of Bitcoin and Other Cryptos," CNBC, December 22, 2024, accessed February 28, 2026, https://www.cnbc.com/2024/12/22/what-google-quantum-chip-breakthrough-means-for-bitcoins-future.html.

15. Mark Webber, Vincent Elfving, Sebastian Weidt, and Winfried K. Hensinger, "The Impact of Hardware Specifications on Reaching Quantum Advantage in the Fault Tolerant Regime," AVS Quantum Science 4, no. 1 (2022): 013801, https://doi.org/10.1116/5.0073075.

16. John Preskill, "Quantum Computing in the NISQ Era and Beyond," Quantum 2 (2018): 79, https://doi.org/10.22331/q-2018-08-06-79.

17. National Institute of Standards and Technology, "NIST Releases First 3 Finalized Post-Quantum Encryption Standards," August 13, 2024, accessed May 21, 2026, https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards.

18. Kris Kwiatkowski and Luke Valenta, "The TLS Post-Quantum Experiment," Cloudflare Blog, October 30, 2019, accessed May 21, 2026, https://blog.cloudflare.com/the-tls-post-quantum-experiment.

19. Apple Security Engineering and Architecture, "iMessage with PQ3: The New State of the Art in Quantum-Secure Messaging at Scale," Apple Security Research Blog, February 21, 2024, accessed May 21, 2026, https://security.apple.com/blog/imessage-pq3.

20. Michael Osborne, Katia Moskvitch, and Jennifer Janechek, "NIST's Post-Quantum Cryptography Standards Are Here," IBM Research Blog, August 13, 2024, accessed May 21, 2026, https://research.ibm.com/blog/nist-pqc-standards.

21. Zoom, "Zoom Bolsters Security Offering with the Inclusion of Post-Quantum End-to-End Encryption in Zoom Workplace," Zoom News, May 21, 2024, accessed May 21, 2026, https://news.zoom.com/post-quantum-e2ee.

22. "Bitcoin Development Mailing List," Google Groups, accessed May 22, 2026, https://groups.google.com/g/bitcoindev.

23. BIP-141 defines witness data as a separate data structure used for transaction validation and moves signatures and script data into that structure. Witness data is organized into a dedicated witness Merkle tree and committed to the block through a witness commitment contained in the coinbase transaction. Since the activation of SegWit, block capacity has been measured using block weight. Non-witness data is counted at four weight units (WU) per byte, while witness

data is counted at one WU per byte. The maximum block weight is 4,000,000 WU. Eric Lombrozo, Johnson Lau, and Pieter Wuille, "BIP 141: Segregated Witness (Consensus Layer)," Bitcoin Improvement Proposals, GitHub, assigned December 21, 2015, accessed May 22, 2026, https://github.com/bitcoin/bips/blob/master/bip-0141.mediawiki.

24. SegWit also mitigated the problem of transaction malleability by changing the way signature data affects transaction identifiers. Transaction malleability refers to the ability to alter a transaction's identifier (transaction ID, or txid) without changing the substance of the transaction itself, typically by modifying signature-related data. Because changes to transaction identifiers can disrupt dependent unconfirmed transactions and payment verification processes, mitigating this issue laid an important foundation for later scalability technologies.

25. The Taproot upgrade is specified across BIP-340, BIP-341, and BIP-342. BIP-340 standardizes 64-byte Schnorr signatures on secp256k1, BIP-341 defines Taproot as a SegWit version 1 output type, and BIP-342 specifies the validation rules for Tapscript, which is used in Taproot spending paths. Pieter Wuille, Jonas Nick, and Tim Ruffing, "BIP 340: Schnorr Signatures for secp256k1," Bitcoin Improvement Proposals, GitHub, assigned January 19, 2020, accessed May 22, 2026, https://github.com/bitcoin/bips/blob/master/bip-0340.mediawiki. Pieter Wuille, Jonas Nick, and Anthony Towns, "BIP 341: Taproot: SegWit Version 1 Spending Rules," Bitcoin Improvement Proposals, GitHub, assigned January 19, 2020, accessed May 22, 2026, https://github.com/bitcoin/bips/blob/master/bip-0341.mediawiki. Pieter Wuille, Jonas Nick, and Anthony Towns, "BIP 342: Validation of Taproot Scripts," Bitcoin Improvement Proposals, GitHub, assigned January 19, 2020, accessed May 22, 2026, https://github.com/bitcoin/bips/blob/master/bip-0342.mediawiki.

26. Hunter Beast, Ethan Heilman, and Isabel Foxen Duke, "BIP 360: Pay-to-Merkle-Root (P2MR)," Bitcoin Improvement Proposals, GitHub, assigned December 18, 2024, accessed May 22, 2026, https://github.com/bitcoin/bips/blob/master/bip-0360.mediawiki.

27. The increase in data size associated with P2MR arises primarily when compared with Taproot's simplest key-path spending method. By contrast, when comparing equivalent script-path spends, P2MR can be smaller than a comparable Taproot script-path spend because it does not include an internal public key. In other words, the cost increase of P2MR depends on the point of comparison. The key trade-off is that reducing the long-term exposure of public keys may increase block-space consumption for certain spending methods.

28. I. Stewart, D. Ilie, A. Zamyatin, S. Werner, M. F. Torshizi, and W. J. Knottenbelt, "Committing to Quantum Resistance: A Slow Defence for Bitcoin against a Fast Quantum Computing Attack," Cryptology ePrint Archive, Paper 2018/213, 2018, accessed May 22, 2026, https://eprint.iacr.org/2018/213.

29. This figure is an estimate derived from block-pattern analysis rather than an established fact. There is no definitive way to determine either Satoshi Nakamoto's exact holdings or the actual ownership status of those coins.

30. BitMEX Research, "Mitigating The Impact of The Quantum Freeze," BitMEX Blog, February 9, 2026, accessed May 22, 2026, https://www.bitmex.com/blog/Mitigating-The-Impact-Of-The-Quantum-Freeze.

31. Marek Palatinus, Pavol Rusnak, Aaron Voisine, and Sean Bowe, "BIP 39: Mnemonic Code for Generating Deterministic Keys," Bitcoin Improvement Proposals, GitHub, assigned September 10, 2013, accessed May 22, 2026,

https://github.com/bitcoin/bips/blob/master/bip-0039.mediawiki. BIP-39 describes a method for implementing mnemonic codes, or mnemonic sentences, for generating deterministic wallets. Under this approach, a sequence of human-readable words is used to generate the wallet's seed, from which cryptographic keys can subsequently be derived.

32. Ethereum Foundation, "Hard Fork Completed," Ethereum Foundation Blog, July 20, 2016, accessed May 23, 2026, https://blog.ethereum.org/2016/07/20/hard-fork-completed.

33. iShares Bitcoin Trust ETF, Prospectus, July 31, 2025, as supplemented November 21, 2025, accessed May 23, 2026, https://www.ishares.com/us/literature/prospectus/p-ishares-bitcoin-trust-12-31.pdf.

Can Bitcoin Survive the Age of Quantum Computing?

The Quantum Threat and the Evolution of Bitcoin Security

Researchers

Yonkyung Kim

Hyemin Son

Corresponding Author

Taemin Oh

Editorial Designer

Seongah Youn

Translator

Samuel J. Hahn

Contact

Hyemin Son hyeomin0109@gmail.com