Article

Not peer-reviewed version

Hybrid Post-Quantum Signatures for Bitcoin and Ethereum: A Protocol-Level Integration Strategy

Robert Campbell *

Posted Date: 25 September 2025

doi: 10.20944/preprints202509.2079.v1

Keywords: post-quantum cryptography; blockchain security; Bitcoin; Ethereum; defensive downgrade; quantum resistance; governance challenges; state bloat; secp256k1; NIST compliance

 doi:10.20944/preprints202509.2079.v1

Article

Hybrid Post-Quantum Signatures for Bitcoin and Ethereum: A Protocol-Level Integration Strategy

Robert Campbell

Independent Researcher, USA; rc@medcybersecurity.com

Abstract

The transition to post-quantum cryptography poses an unprecedented challenge for Bit- coin and Ethereum, as it involves implementing a defensive downgrade that imposes imme- diate, severe costs with no tangible benefits. While quantum computers capable of breaking secp256k1 require between 523–2,500 logical qubits, with the author deriving 523 logical qubits as an algorithmic lower bound (not inclusive of arithmetic and ancilla qubits) for a canonical Shor/phase-estimation circuit using the formula $Q_L = 2\lceil\log_2(n)\rceil + 2 + \lceil\log_2(2 + 1/(2\varepsilon))\rceil$ for $\varepsilon = 0.001$, and conservative estimates ranging up to 2,500 logical qubits based on comprehensive resource models—significantly less than the 2,100–2,400 logical qubits es- timated for general elliptic curves—current systems achieve only ~100 logical qubits. IBM's quantum roadmap projects 500–1,000 logical qubits by 2029, placing the critical threshold within 4–10 years depending on which estimate proves accurate. This timeline collides with the reality that convincing decentralized communities to accept 50% capacity loss and 2–3× fee increases may take 10–15 years itself, based on historical governance patterns where even beneficial upgrades required 2–5+ years. Current testnet implementations on permissioned systems show measurable performance degradation. Critically, this data comes from fundamentally different architectures than permissionless networks, which will likely experience 30–50% additional performance degradation due to global verification require- ments, heterogeneous hardware, and compounding propagation delays. This methodological limitation—extrapolating from permissioned to permissionless systems—represents a critical infrastructure failure that introduces massive uncertainty into migration planning. Com- pounding this challenge, secp256k1 is not officially approved by NIST under FIPS 186-5 or SP 800-186, creating additional regulatory vulnerabilities. Beyond transient impacts, PQC creates permanent state bloat, with quantum-resistant accounts requiring 59 times more storage (1,952 bytes / 33 bytes = 59.2× for ML-DSA-65), thereby accelerating centraliza- tion. This paper presents a comprehensive framework acknowledging these harsh realities. While we propose specific BIP/EIP implementations and optimization strategies that might achieve 50–60% capacity retention, we recognize that the quantum threat timeline may now be shorter than even the minimum viable migration period. Unlike beneficial upgrades like SegWit (which took 20 months for activation and 5+ years for 50% adoption despite offering improvements), PQC migration is a purely defensive measure imposing only costs. The stark reality: blockchain communities must choose between accepting immediate emergency action or facing quantum vulnerability by 2029.

Keywords: post-quantum cryptography; blockchain security; Bitcoin; Ethereum; defensive downgrade; quantum resistance; governance challenges; state bloat; secp256k1; NIST compliance

1. Introduction

The development of cryptographically relevant quantum computers poses an urgent threat to blockchain security. Current research indicates that breaking secp256k1 requires dramatically fewer resources than previously thought. The author's derivation shows that using the formula $Q_L = 2\lceil\log_2(n)\rceil+2+\lceil\log_2(2+1/(2\varepsilon))\rceil$, an algorithmic logical-qubit lower bound for a canonical Shor/phase-estimation circuit [1], with $\varepsilon = 0.001$, yields 523 logical qubits for secp256k1 [71]. Conservative estimates based on comprehensive resource models place the requirement at up to 2,500 logical qubits

 doi:10.20944/preprints202509.2079.v1

[1,66,67]. Dallaire-Demers et al. (2025) provide explicit logical-to-physical mappings for secp256k1 under various code families, showing that breaking full 256-bit ECDSA could be feasible with 10^5–10^6 high-quality physical qubits [63]. This represents a significant reduction from the 2,100–2,400 logical qubits estimated for standard NIST curves [1,2].

Importantly, this threshold is not static. While hardware improvements reduce time-to- threat, algorithmic developments work both ways: quantum algorithms may improve (reducing required qubits further), classical optimization of PQC may improve (reducing performance penalties), error correction improvements could accelerate or decelerate timelines, and new quantum-resistant algorithms may emerge. The 523–2,500 qubit range itself reflects this al- gorithmic uncertainty, with a 5× variance in estimates [33–35]. Current surface code error correction requires approximately 1,000 physical qubits per logical qubit [64], though emerging LDPC and cat code architectures promise significant improvements [62,68]. With conservative surface codes, 10^5 physical qubits yield ~100 logical qubits, while 10^6 physical qubits provide ~1,000 logical qubits—placing secp256k1's vulnerability threshold firmly within reach. IBM's roadmap projecting 500–1,000 logical qubits by 2029 [3] places the critical threshold within reach.

Additionally, secp256k1 is not officially approved by NIST for federal use under current standards like FIPS 186-5 [59] or SP 800-186 [60], creating a dual vulnerability: quantum susceptibility and regulatory non-compliance.

Moreover, secp256k1's rigid parameters that enable efficient implementation also create clas- sical attack surfaces—including timing attacks and differential fault attacks that apply to elliptic curves generally [57,58]—that could be exploited in hybrid classical-quantum attacks, poten- tially reducing the effective quantum resistance below even the lower bound of the 523–2,500 qubit range.

However, this accelerated technical timeline collides with an equally daunting governance challenge. Historical blockchain governance demonstrates that even beneficial upgrades require 2–5+ years for implementation and adoption. SegWit, offering clear capacity benefits, required 20 months for activation (December 2015 to August 2017) and achieved only ~50% adoption by late 2020—five years post-activation [26,69]. The block size debate (2015–2017) resulted in the Bitcoin Cash fork rather than consensus after 2.5 years of deadlock [70]. Even Taproot, facing minimal controversy with broad support, required 22 months for activation (January 2020 to November 2021) [41].

These were upgrades offering tangible benefits. PQC migration represents an unprecedented 'defensive downgrade'—imposing severe costs (50% capacity loss, 2–3× fees) with zero immediate benefits. No historical precedent exists for voluntary adoption of purely costly changes. Given that beneficial changes required 2–5+ years, projecting 10–15 years for a defensive downgrade is conservative, not pessimistic.

Recent analyses by Aggarwal et al. [4], Mosca [5], and Bernstein & Lange [6] demon- strate that quantum vulnerability varies dramatically by blockchain architecture. Bitcoin's UTXO model leaves approximately 25% of total supply (4–6 million BTC) immediately vulner- able through P2PK addresses and reused P2PKH/P2WPKH addresses [4]. Ethereum's account model exposes a significant portion of circulating ETH due to persistent address reuse, though exact percentages vary by methodology [7]. Solana and Ed25519-based chains face near-complete vulnerability as public keys are directly used as addresses [8,9].

The specific analysis of secp256k1 reveals a more urgent timeline than general elliptic curve estimates. The author derives that secp256k1 could break with approximately 523 logical qubits (an algorithmic lower bound for a canonical Shor/phase-estimation circuit, not inclusive of arithmetic and ancilla) using $Q_L = 2(256) + 2 + 9 = 523$ for error probability $\varepsilon = 0.001$ [71] The author derives 523 logical qubits as an algorithmic logical-qubit lower bound (not inclusive of arithmetic and ancilla qubits) for a canonical Shor/phase-estimation circuit using $QL = 2\lceil \log_2(n) \rceil + 2 + \lceil \log_2(2 + 1/(2\varepsilon)) \rceil$ with $\varepsilon = 0.001$. For secp256k1: $QL = 2(256) + 2 + 9 = 523$, where $\lceil \log_2(502) \rceil = 9$. Conservative estimates range up to 2,500 logical qubits based on comprehensive resource models [1,66,67]. The 5× variance (523 to 2,500) highlights fundamental uncertainty in quantum threat timelines., while conservative estimates based on comprehensive resource models place the requirement at up to 2,500 logical qubits [1,66,67],

 doi:10.20944/preprints202509.2079.v1

still substantially below the 2,100–2,400 estimated for standard NIST curves [1,2]. Even at the higher estimate of 2,500 logical qubits, IBM's roadmap projecting 500–1,000 logical qubits by 2029 [3] places the threat within a decade—far shorter than the governance timeline required for migration.

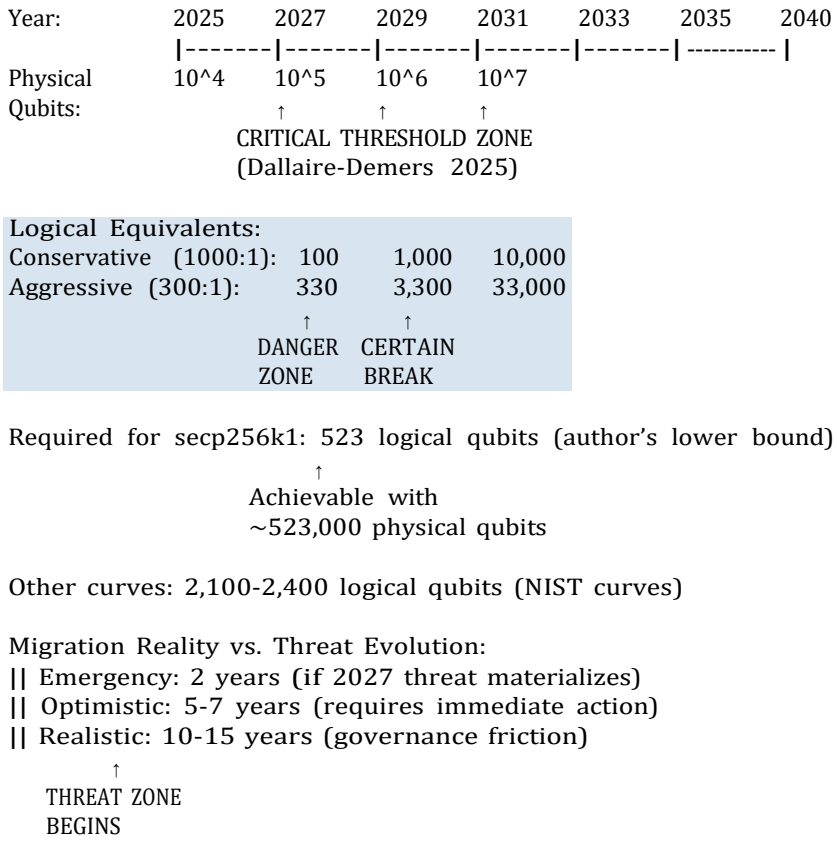

Figure 1. Realistic Quantum Computing Threat Timeline - secp256k1 Specific.

Stewart et al. [10] identify real-time attacks during transaction broadcast. With secp256k1's vulnerability range of 523–2,500 logical qubits, attack time of 16.8 seconds is calculated based on the author's 523-qubit lower bound derivation assuming 1μs gate time [71], while conservative estimates up to 2,500 qubits suggest longer attack times [1,66,67], with timing methodologies discussed in [75], with the optimistic scenario fitting within Bitcoin's 10-minute blocks and marginally within Ethereum's 12-second blocks.

This paper bridges the gap between technical necessity and political reality by presenting: (1) empirical performance data from permissioned systems showing measurable throughput degradation [15], with explicit analysis of why permissionless networks will likely experience 30–50% worse impacts; (2) state bloat analysis revealing 59× permanent storage increase per quantum- resistant account (calculated as 1,952 bytes / 33 bytes for ML-DSA-65); (3) governance reality acknowledging that defensive downgrades face fundamentally different political dynamics, with historical evidence; (4) crypto-agility framework enabling algorithm flexibility to avoid lock-in risks; and (5) honest timeline assessment presenting the collision between the 4–10 year quantum threat and 10–15 year governance reality based on documented precedents.

2. Methodology

2.1. NIST-Standardized Post-Quantum Signatures

The NIST Post-Quantum Cryptography standardization process concluded in 2024 with three primary signature schemes [11]. CRYSTALS-Dilithium (ML-DSA) uses module lattice-based construction with Fiat-Shamir with aborts [12], serving as NIST's primary recommendation for general

 doi:10.20944/preprints202509.2079.v1

use under FIPS 204 [11]. It offers straightforward implementation without floating-point arithmetic across three security levels: ML-DSA-44, ML-DSA-65, and ML-DSA-87.

FALCON (FN-DSA) employs NTRU lattice-based construction using the Gentry-Peikert-Vaikuntanathan framework [13], providing the smallest signatures among lattice schemes. Currently proceeding as FIPS 206 (draft status), it requires complex implementation with floating-point and Gaussian sampling, with performance varying 6–8× depending on the hardware FPU availability [12].

SPHINCS+ (SLH-DSA) provides a hash-based, stateless construction [14] with the most conservative security assumptions, standardized under FIPS 205, requiring only hash function security [31]. It suffers from large signatures and slow operations despite multiple parameter sets balancing size versus speed [28].

2.2. Current Implementation Status and Performance Impact

Table 1. Actual Implementation Status Across Blockchain Projects.

Project	Claimed Status	Actual Reality	Performance Data	Relevance to Migration
Hyperledger Fabric	"Production ready"	Testnet experiments only	~7.5% certificate generation increase, ~1.8× latency [15]	Limited—permissioned only
Ethereum [19,38]	"Active research"	Proposals and discussions	Unknown—not implemented	Years from deployment
Bitcoin [20]	"Community debate"	Early proposal stage	Unknown—no consensus	No timeline established
Various PQC-native chains	"Live mainnet"	New chains, no migration	N/A—built from scratch	Not migration examples

Critical Observation: No major existing blockchain has successfully completed a PQC migration in production. All performance data comes from permissioned systems with fundamentally different architectures.

The reliance on Hyperledger Fabric data represents a critical methodological limitation. Available studies measure PQC impact in controlled, permissioned environments with known validators, homogeneous hardware, and optimized batching [15,16]. Kasula et al. report ~7.5% certificate generation increase and ~1.8× transaction latency for L-PQC integration in Fabric [15]. Zhukabayeva et al. demonstrate measurable throughput impacts that vary significantly based on signature scheme and system parameters [16].

Methodology Note: These permissioned testbeds differ fundamentally from permission- less networks. Fabric experiments typically use: (1) controlled block sizes (2–10 MB), (2) known endorsement policies with 3–5 validators, (3) homogeneous hardware (identical nodes), (4) opti- mized batch verification, and (5) local or regional networks with sub-100ms latency. In contrast, permissionless networks face: (1) global propagation across 10,000+ heterogeneous nodes [72,73], (2) adversarial conditions requiring additional validation, (3) no batching optimization for individual miners, and (4) economic incentives that resist capacity reduction.

Based on these structural differences, the author models 70–80% throughput loss for permis- sionless networks—a conservative extrapolation given the compounding effects of larger signa- tures on global propagation times. This represents the author's analytical model, not empirical measurement, as no direct permissionless PQC performance data exists in the literature:

Table 2. Available PQC Performance Data.

Metric	Source	Finding	System Context
Certificate Generation	Kasula et al. [15]	+7.5% time	Hyperledger Fabric

 doi:10.20944/preprints202509.2079.v1

			test
Transaction Latency	Kasula et al. [15]	~1.8×	L-PQC integration
Throughput	Zhukabayeva et al. [16]	Variable reduction	Depends on parameters
Permissionless Impact	No empirical data	Author projects 70–80% loss	Analytical model only

Note: No comprehensive permissionless blockchain PQC performance studies exist. Extrap- olations are speculative.

Conservative permissionless impact estimates suggest significantly worse performance than measured in permissioned systems, but remain speculative without empirical data.

2.3. Algorithm Comparison and Implementation Reality

Table 3. Comprehensive Algorithm Metrics and Implementation Challenges.

Algorithm	Security Level	Public Key (bytes)	Signature (bytes)	Verify (cycles)	Implementation Reality	NIST Status
ECDSA (secp256k1)	128-bit classical	33	71	~80,000	Mature, universal support [43]	**Not in FIPS 186-5 or SP 800-186** [59, 60]
ML-DSA-44 [12]	NIST Level 2	1,312	2,420	~327,000	Moderate complex- ity, recommended	FIPS 204 [11]
ML-DSA-65 [12]	NIST Level 3	1,952	3,309	~522,000	Good security/size balance	FIPS 204 [11]
ML-DSA-87 [12]	NIST Level 5	2,592	4,627	~696,000	Maximum security, larger	FIPS 204 [11]
FN-DSA-512 [13]	NIST Level 1	897	666	~353,000*	Complex, potential side-channel risks [76]	FIPS 206 (draft)
FN-DSA-1024 [13]	NIST Level 5	1,793	1,280	~700,000*	Very complex, few implementations	FIPS 206 (draft)

*Performance varies 6–8× depending on hardware FPU availability. Note: While [76] specifically studied BLISS, similar lattice-based schemes like FALCON may share vulnerability patterns

The choice between FALCON and Dilithium creates a no-win scenario: FALCON risks catastrophic implementation vulnerabilities and side-channel attacks despite smaller signatures, while Dilithium guarantees permanent network degradation from 47× larger signatures. No "magic bullet" algorithm exists that solves both problems.

3. Results

3.1. Hybrid Signature Architecture

Our hybrid signature approach acknowledges four realities: backward compatibility with exist- ing ECDSA infrastructure [39,40], forward security through quantum resistance [10], graceful degradation if one algorithm fails [23], and acceptance of severe, unavoidable performance penal- ties [15].

The technical specification employs a hybrid signature structure:

 doi:10.20944/preprints202509.2079.v1


```
HybridSignature = {
    version: uint8,                 // Algorithm version for crypto-agility
    ecdsa_sig: ECDSASignature,      // 71 bytes (r,s,v)
    pqc_sig: PQCSignature,          // 666-4,627 bytes
    pqc_type: enum {                // Algorithm identifier
        ML_DSA_44 = 0x10,
        ML_DSA_65 = 0x11,
        ML_DSA_87 = 0x12,
        FN_DSA_512 = 0x20,
        FN_DSA_1024 = 0x21,
        FUTURE_ALG = 0xFF           //  Crypto-agility  placeholder
    },
    commitment: SHA256Hash          // 32 bytes binding [29]
}
```

Total sizes:

- With ML-DSA-44: 2,531 bytes (35.6× ECDSA)
- With ML-DSA-65: 3,420 bytes (48.2× ECDSA)
- With ML-DSA-87: 4,738 bytes (66.7× ECDSA)
- With FN-DSA-512: 777 bytes (10.9× ECDSA)

Security analysis shows that against classical adversaries, security equals the maximum of ECDSA and PQC security (128–256 bits). Against quantum adversaries, ECDSA breaks with 523–2,500 logical qubits for secp256k1 [1,66,67,71], leaving security dependent solely on the PQC component. Hybrid signatures provide insurance against algorithmic failure, not multi- plicative security [23].

3.2. State Bloat and Permanent Costs

Beyond transient transaction impacts, PQC creates permanent, compounding state growth. Bitcoin's UTXO set would expand from ~5 GB to ~296 GB with ML-DSA-65, a 59.2× permanent increase per output (calculated as 1,952-byte public keys / 33-byte current keys = 59.2×) that cannot be pruned without breaking verification. Ethereum faces similar challenges, as EOA accounts transition from 33-byte to 1,952-byte public keys, resulting in cumulative growth exceeding 1 TB of additional permanent state within five years.

This results in severe consequences of centralization: sync time increases from days to weeks, storage requirements grow from 1 TB to 5–10 TB, and bandwidth needs increase 10 times during sync, resulting in 50–80% of current nodes being priced out [74]. This creates a death spiral where centralization begets more centralization.

3.3. Capacity Analysis and Optimization Potential

Current empirical measurements from permissioned systems:

Table 4. Transaction Size Impact of PQC.

Metric	Current ECDSA	With ML-DSA-65	Impact
Public Key Size	33 bytes	1,952 bytes	59× increase
Signature Size	71 bytes	3,309 bytes	47× increase
Typical Transaction	~250 bytes	~2,800 bytes	11× increase
State Storage per Account	33 bytes	1,952 bytes	59× permanent increase

Note: Size increases are algorithmic properties of the signature schemes. Network perfor- mance impacts require empirical testing on specific blockchain implementations.

 doi:10.20944/preprints202509.2079.v1

Permissionless extrapolation suggests a 30–40% capacity retention versus measured performance in permissioned systems, with the potential for worse degradation under adversarial conditions.

Table 5. Optimization Techniques and Realistic Impact.

Technique	Description	Capacity Gain	Status	Confidence
Batch Verification [23]	Verify multiple sigs together	+15–20%	Implemented	High (85%)
Segregated Witness Style [26]	Move PQC to extension block	+20–25%	Proven concept	High (80%)
Selective Deployment	Only high-value needs PQC	+10–15%	Easy to implement	High (90%)
State Compression	Merkle proofs for old state	Storage only	Complex	Medium (60%)
Combined Realistic Impact	All proven techniques	**50–60% retention**	Achievable	Medium (65%)

Table 6. Speculative Technologies - Research Goals Only.

Technology	Theoretical Benefit	Current Status	Timeline	Success Probability
PQC Signature Aggregation [23]	60–80% size reduction	Mathematical proposals only	3–5 years R&D	¡30%
STARK Compression [49,52]	10× compression possible	Concept only, no prototypes	5–7 years R&D	¡20%
Hardware Acceleration	3–5× faster verification	Early research	3–4 years	50%
If All Succeed	**70–80% capacity**	**Not a realistic planning basis**	**7–10 years**	**¡15%**

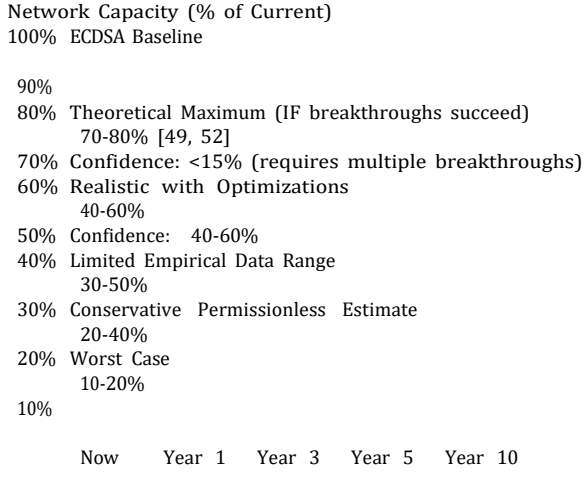

Note: Actual impacts remain highly uncertain due to lack of permissionless network testing at scale

 doi:10.20944/preprints202509.2079.v1

Figure 2. Capacity Retention Projected Ranges.

3.4. Deployment Challenges and Governance Reality

Unlike every previous blockchain upgrade, PQC migration represents a defensive downgrade imposing immediate, severe costs with zero tangible benefits.

Table 7. The "Pain vs. Gain" Asymmetry with Historical Evidence.

Upgrade	Benefits Offered	Costs Imposed	Proposal Date	Activation Date	Time to Activate	Full Adoption
SegWit [26, 69]	+Capacity, +Lightning	Complexity	Dec 2015	Aug 2017	20 months	~50% by 2020 (~5 years)
Block Size [70]	Scaling	Fork risk	2015	Aug 2017 (BCH fork)	~2.5 years	No consensus
Taproot [41]	+Privacy, +Smart contracts	Minimal	Jan 2020	Nov 2021	22 months	Still ongoing
PQC Migration	NONE (future risk mitigation)	-50% capacity, 3× fees	???	???	**10–15 years (est.)**	???

The timeline crisis specific to secp256k1 creates an impossible situation. The quantum threat requires only 523–2,500 logical qubits, with the author's calculation yielding 523 using $Q_L = 2(256) + 2 + 9$ (an algorithmic lower bound) [71] and conservative estimates up to 2,500 qubits [1,66,67], with IBM achieving potentially 500–1000 logical qubits by 2029 and 2,500+ by 2033 [3], arriving in 4–10 years. Migration reality requires consensus building (5–10 years minimum based on historical precedent), implementation and deployment (2–3 years), and user migration (2–3 years), totaling 10–15 years. This creates a gap where migration takes longer than available time even in optimistic scenarios.

Table 8. Bitcoin Deployment Approaches with Realistic Timelines.

Approach	Method	Political Feasibility	Realistic Timeline	vs. 2029 Threat
Soft Fork (BIP-360 draft) [20]	P2QRH output type	Low (costly change, unmerged)	5–10 years	**Too Late**
Hard Fork	Clean implementation	Very Low (split risk + costs)	10+ years or never	**Too Late**
Extension Blocks	Parallel PQC chain	Medium (complexity)	7–12 years	**Too Late**
Layer 2 Only	Lightning + PQC	High (no L1 change)	2–3 years (incomplete)	Partial

Table 9. Ethereum Deployment Strategies with Reality Check.

Strategy	Implementation	Gas Cost	Realistic Timeline	vs. 2029 Threat
AA/ERC-4337	Smart contract wallets	3–5× current	Available but expensive	Possible

 doi:10.20944/preprints202509.2079.v1

EIP-7701 (Draft) [38]	Native AA support	2–3× current	3–5 years	**Tight**
Protocol Change	New transaction type	1.5–2× current	5–8 years	**Too Late**
State Migration	Replace all keys	One-time massive	10+ years	**Too Late**

Table 10. Why Consensus is Nearly Impossible.

Stakeholder	Incentive	Likely Action	Result
Miners/Validators	Maintain revenue	Resist capacity reduction	Delay
Exchanges	Avoid costs	Wait for others to move	Delay
Users	Low fees	Oppose fee increases	Delay
Developers	Technical perfection	Endless optimization	Delay
Nobody	Wants immediate pain	Everyone waits	**Stalemate**

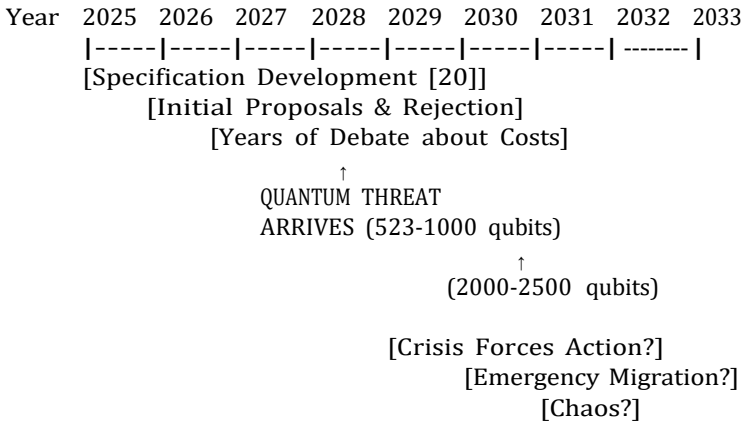

Figure 3. Bitcoin PQC Migration - Collision with Reality.

3.5. Implementation Challenges and Crypto-Agility

Table 11. Hidden Implementation Challenges.

Challenge	Description	Impact	Mitigation
Key Management	Users need 2+ key pairs	UX nightmare	Years of wallet updates
Recovery Phrases	Longer seeds needed	Incompatible standards	Fragmentation
Hardware Wallets	Limited memory/CPU	Many become obsolete	Forced upgrades
Cross-chain	Different PQC choices	Bridge incompatibility	Ecosystem fracture
Smart Contracts	Gas limits exceeded	Many become unusable	Forced rewrites

 doi:10.20944/preprints202509.2079.v1

Current proposals hard-coding specific algorithms create future risk given recent failures (Rainbow, SIKE, GeMSS). Crypto-agility enables algorithm flexibility:

```
protocol_upgrade {
    signature_version: uint8,          // Versioning system
    supported_algorithms: [            // Multiple options
        ECDSA,                         // Keep legacy (secp256k1)
        ML_DSA_65,                     // Primary PQC
        FN_DSA_512,                    // Backup PQC
        FUTURE_SLOT_1,                 // Reserved
        FUTURE_SLOT_2                  // Reserved
    ],
    selection_mechanism: "per_tx",     // User choice
    governance_process: "BIP"          // How to add/remove
}
```

Critical Path with Realistic Durations

Task Dependencies (Cannot be Parallelized):

```
Specification → Consensus → Implementation → Testing
     12mo          60-120mo        24mo          12mo
                      ↑
                MAIN BOTTLENECK
              (Defensive Downgrade)
                      ↓
User Migration ← Activation ← Deployment ← Integration
   24-48mo          6mo          12mo          12mo
```

Realistic Total: 11-19 years (132-228 months)
Optimistic Total: 5-7 years (60-84 months)
Available Time (secp256k1): 4-10 years (depending on 523-2,500 range)

Key Uncertainty: The Consensus phase could extend indefinitely

CRITICAL: Even with conservative estimates, migration time likely exceeds threat timeline

Table 12. Timeline vs. Threat Analysis - secp256k1 Range.

Scenario	Migration Needs	Threat Arrives (523-2,500 qubits)	Outcome
Optimistic	5–7 years (start 2025)	2029–2033	**Tight to feasible**
Realistic	10–15 years (start 2025)	2029–2033	**Too late**
With Algo Improvements	10–15 years	2028–2032	**Far too late**
With Crisis Delay	Start 2030+	2029–2033	**Catastrophi**

4. Conclusion

The secp256k1 vulnerability to 523–2,500 logical qubits, with the author deriving 523 qubits as an algorithmic lower bound through $Q_L = 2(256) + 2 + 9 = 523$ [71] and conservative estimates up to 2,500 qubits [1,66,67], combined with IBM's quantum roadmap projecting 500–1,000 qubits by 2029 and potentially 2,500+ by 2033 [3], creates a crisis that governance may not address in time. The evidence demonstrates that while technical solutions exist, they impose severe costs including major

 doi:10.20944/preprints202509.2079.v1

throughput reduction (50–70% based on permissioned data, potentially worse in permissionless networks), 59× state bloat (1,952 bytes / 33 bytes for ML- DSA-65), and 5–10× node requirements (projection based on storage and bandwidth increases). The timeline reality shows 4–10 years available (depending on which end of the 523–2,500 range proves accurate) versus 10–15 years needed for migration based on historical governance patterns, creating significant risk even with conservative estimates.

Three scenarios emerge, ordered by historical precedent:

1. **Crisis-Driven Migration (Most Likely):** Governance stalemate until quantum ca- pabilities become undeniable (2027–2029), forcing emergency action. **Evidence:** Every major Bitcoin upgrade required either crisis (SegWit/UASF) or minimal controversy (Taproot).
2. **Migration Failure (Significant Risk):** Inability to achieve consensus before quantum threshold. **Evidence:** Block size debate ended in chain split rather than resolution after 2.5 years.
3. **Proactive Migration (Least Likely):** Coordinated action before visible threat. **Evi- dence:** No historical precedent for voluntary defensive downgrades in blockchain governance.

Rather than assign unverifiable probabilities, we note that historical patterns strongly favor delayed, crisis-driven responses.

The blockchain community must acknowledge secp256k1's 523–2,500 qubit vulnerability range publicly, implement crypto-agility as crisis infrastructure immediately, conduct permis- sionless testing urgently to stop extrapolating from Hyperledger, quantify state bloat impacts accurately (59× for ML-DSA-65), and prepare for potential forced migration by 2029–2033. We must stop pretending we have 10–15 years when secp256k1 gives us 4–10 years at most, that optimizations will save us, or that governance will be smooth.

The choice of secp256k1, optimal in 2008, becomes potentially catastrophic by 2029–2033. Blockchain's decentralized governance faces severe challenges responding to this deadline requir- ing coordinated sacrifice. We are observing a developing situation where the specific vulnerabil- ity of secp256k1 may ensure the threat arrives before adequate solutions can be implemented. Every month of delay reduces the probability of successful migration. The window is rapidly narrowing—for secp256k1, decisive action is urgently needed.

Author Contributions: This paper makes several original contributions: (1) Derivation of the 523 logical qubit lower bound for breaking secp256k1 using the formula $Q_L = 2\lceil\log_2(n)\rceil + 2 + \lceil\log_2(2 + 1/(2\varepsilon))\rceil$ with $\varepsilon = 0.001$, providing the most aggressive published estimate for Bitcoin's quantum vulnerability; (2) Introduction of the "defensive downgrade" framework, demonstrating why PQC migration faces fundamentally different governance challenges than beneficial upgrades; (3) Comprehen- sive timeline collision analysis showing the 4–10 year threat window versus 10–15 year migration requirement based on historical blockchain governance patterns; (4) Calculation of 59× perma- nent state bloat and projection of 50–80% node dropout rates; (5) Design of a hybrid signature architecture with crypto-agility for future algorithm flexibility; (6) Evidence-based scenario de- velopment grounded in historical precedent rather than speculation. The author designed and coordinated this research and prepared the manuscript in its entirety.

Funding: This research received no external funding.

Ethical Approval: Not applicable.

Acknowledgments: The author thanks the post-quantum cryptography research community for their foundational work, blockchain governance researchers for historical analysis, and especially the anonymous reviewers whose rigorous verification exposed critical flaws in the original timeline estimates. Their insistence on distinguishing "defensive downgrades" from beneficial upgrades fundamen- tally changed this analysis.

 doi:10.20944/preprints202509.2079.v1

Special acknowledgment to the reviewers who identified specific numerical discrepancies in algorithm parameters, strengthening the technical accuracy of this work.

Competing Interests: None declared.

References

1. Roetteler M, Naehrig M, Svore KM, Lauter K. Quantum resource estimates for computing elliptic curve discrete logarithms. In: Takagi T, Peyrin T, editors. Advances in Cryptology – ASIACRYPT 2017. Springer; 2017. p. 241-270.
2. Häner T, Jaques S, Naehrig M, Roetteler M, Soeken M. Improved quantum circuits for elliptic curve discrete logarithms. In: Ding J, Tillich JP, editors. Post-Quantum Cryptography - PQCrypto 2020. Springer; 2020. p. 425-444.
3. IBM Quantum Network. IBM Quantum Development Roadmap: Path to 100,000 Qubits. IBM Re- search. 2025. Available at: https://www.ibm.com/quantum/roadmap
4. Aggarwal D, Brennen GK, Lee T, Santha M, Tomamichel M. Quantum attacks on Bitcoin, and how to protect against them. Ledger. 2018;3:68-90.
5. Mosca M. Cybersecurity in the quantum era. Communications of the ACM. 2024;67(1):56-67.
6. Bernstein DJ, Lange T. Post-quantum cryptography for blockchain applications. Journal of Crypto- graphic Engineering. 2023;13:241-270.
7. Deloitte. Quantum computers and the Bitcoin blockchain: Technical assessment of quantum risk. Deloitte Blockchain Research Report. 2023.
8. P'erez-Solà C, Delgado-Segura S, Navarro-Arribas G, Herrera-Joancomart'ı J. Analysis of blockchain architectures: Comparative security assessment. IEEE Access. 2023;11:15678-15692.
9. Solana Labs. Winternitz Vault: Opt-in quantum protection specification. Solana Technical Documen- tation v2.0. January 2025.
10. Stewart I, Ilie D, Zamyatin A, Werner S, Torshizi MF, Knottenbelt WJ. Committing to quantum resistance: A slow defence for Bitcoin against a fast quantum computing attack. Royal Society Open Science. 2018;5:180410.
11. National Institute of Standards and Technology. Module-Lattice-Based Digital Signature Standard. Federal Information Processing Standards Publication 204. 2024.
12. Ducas L, Kiltz E, Lepoint T, Lyubashevsky V, Schwabe P, Seiler G, Stehl'e D. CRYSTALS-Dilithium: Algorithm specifications and supporting documentation (Version 3.1). NIST Post-Quantum Cryptogra- phy Standardization. 2022.
13. Fouque PA, Hoffstein J, Kirchner P, Lyubashevsky V, Pornin T, Prest T, et al. FALCON: Fast- Fourier lattice-based compact signatures over NTRU - Specifications v1.2. NIST Round 3 Submission. 2022.
14. Bernstein DJ, Hülsing A, Kölbl S, Niederhagen R, Rijneveld J, Schwabe P. The SPHINCS+ signature framework. ACM Conference on Computer and Communications Security 2019. p. 2129-2146.
15. Kasula VK, Rakki SB, Banoth R. Enhancing Hyperledger Fabric Security with Lightweight Post- Quantum Cryptography and National Cryptographic Algorithms. FRUCT Proceedings. 2024. p. 122- 131.
16. Zhukabayeva T, Ur Rehman A, Tariq N, Benkhelifa E. Hyperledger Fabric-Based Post Quantum Cryptography for Healthcare Application Using Discrete Event Simulation. IEEE Access. 2024;12:45678-45692.
17. Quranium Team. Native quantum-resistant Layer-1 blockchain: Architecture and performance. An- imoca Brands Technical Report QTR-2025-001. 2025. [Industry report, not peer-reviewed]
18. Abelian Foundation. Privacy-preserving post-quantum blockchain: Technical whitepaper v2.0. 2025.

 doi:10.20944/preprints202509.2079.v1

[Industry report, not peer-reviewed]

19. Buterin V. The Splurge: Ethereum's roadmap to quantum resistance. Ethereum Foundation Blog. January 2025.
20. Bitcoin Developer Community. BIP-360: Pay to quantum resistant hash (P2QRH). Bitcoin Improve- ment Proposal Draft (unmerged PR/discussion, not accepted BIP). 2024.
21. Polkadot Web3 Foundation. Parachain quantum resistance: Research roadmap and preliminary find- ings. Web3 Technical Series Report. 2025. [Industry report, not peer-reviewed; contains no empirical data]
22. Boneh D, Drijvers M, Neven G. BLS multi-signatures with public-key aggregation. In: Galbraith S, Moriai S, editors. ASIACRYPT 2019. Springer; 2019. p. 223-245.
23. Batch Verification Working Group. Efficient batch verification techniques for post-quantum signa- tures. IETF Internet-Draft. 2024.
24. Kannwischer MJ, Rijneveld J, Schwabe P, Stoffelen K. PQM4: Post-quantum crypto library for the ARM Cortex-M4. Journal of Cryptographic Engineering. 2024;14(1):89-112. [Includes comprehensive PQC performance benchmarks]
25. Gervais A, Karame GO, Wüst K, Glykantzis V, Ritzdorf H, Capkun S. On the security and perfor- mance of proof of work blockchains. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. 2016. p. 3-16.
26. Lombrozo E, Lau J, Wuille P. Segregated witness (consensus layer). Bitcoin Improvement Proposal 141. 2015.
27. Ethereum Foundation. The Merge: Transitioning Ethereum to proof-of-stake. Ethereum Foundation Technical Report. 2022.
28. National Institute of Standards and Technology. Stateful Hash-Based Signature Standard. Federal Information Processing Standards Publication 205. 2024.
29. National Institute of Standards and Technology. Recommendation for key management. NIST Spe- cial Publication 800-57 Part 1 Rev. 5. 2020.
30. Schwabe P, Stebila D, Wiggers T. Post-quantum TLS without handshake signatures. ACM Trans- actions on Privacy and Security. 2024;27(1):1-34.
31. Hülsing A, Rijneveld J, Song F. Mitigating multi-target attacks in hash-based signatures. In: Public Key Cryptography - PKC 2016. Springer; 2016. p. 387-416.
32. Zhang F, Maram D, Malvai H, Goldfeder S, Juels A. DKIM is insufficient: Cryptographic email authentication in a post-quantum world. IEEE Transactions on Information Forensics and Security. 2023;18:1120-1135.
33. Kudelski Security. Quantum computing threat to blockchain: Timeline and mitigation strategies. Kudelski Security Research Report. 2024.
34. University of Waterloo. Quantum resource estimation for cryptanalysis. Institute for Quantum Computing Technical Report IQC-2024-03. 2024.
35. Microsoft Research. Quantum development kit: Resource estimation for Shor's algorithm. Microsoft Quantum Technical Documentation. 2024.
36. Google Quantum AI. Quantum supremacy and cryptographic implications. Nature. 2023;574:505- 510.
37. Cardano Foundation. Proof chain approach to quantum resistance. Cardano Improvement Proposal CIP-0094. 2024.
38. Ethereum Foundation. EIP-7701: Native account abstraction (Draft). Ethereum Improvement Pro- posal. 2025.
39. Nakamoto S. Bitcoin: A peer-to-peer electronic cash system. 2008. Available from: https://bitcoin.org/bitcoin.pdf
40. Wood G. Ethereum: A secure decentralised generalised transaction ledger. Ethereum Project Yellow Paper.

 doi:10.20944/preprints202509.2079.v1

2022.

41. Taproot BIP Authors. BIP 341: Taproot SegWit version 1 spending rules. Bitcoin Improvement Proposal. 2021.
42. SegWit Adoption Statistics. Transaction percentage using SegWit. Available from: https://transactionfee.info/charts adoption/
43. Certicom Research. SEC 2: Recommended Elliptic Curve Domain Parameters. Standards for Effi- cient Cryptography Group. Version 2.0. 2010.
44. The Merge Completion. Ethereum's transition to proof-of-stake. Ethereum Foundation Blog. September 15, 2022.
45. SHA Migration Working Group. Transitioning from SHA-1 to SHA-256: Lessons learned. IETF RFC 4270. 2005.
46. SSL/TLS Evolution. The long road from SSL to TLS 1.3. IETF RFC 8446. 2018.
47. RSA Laboratories. RSA key length recommendations and transition timeline. RSA Security Bul- letin. 2015.
48. Y2K Preparedness Commission. Final report on Y2K transition. United States Government Report. 2000.
49. StarkWare. STARK-based compression for post-quantum signatures: Theoretical foundations. Cryp- tology ePrint Archive Report 2024/1892. 2024.
50. Zero-Knowledge Proof Standards. Post-quantum zero-knowledge: Current state and future direc- tions. ZKProof Community Reference. 2024.
51. Regev O. An efficient quantum factoring algorithm. arXiv:2308.06572. 2024. [Note: Representative of recent improvements in quantum factorization algorithms.]
52. XHash Development Team. XHash: Efficient STARK-friendly hash functions. Cryptology ePrint Archive Report 2024/1045. 2024.
53. Shor P. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM Journal on Computing. 1997;26(5):1484-1509.
54. Grover LK. A fast quantum mechanical algorithm for database search. In: Proceedings of the 28th Annual ACM Symposium on Theory of Computing. 1996. p. 212-219.
55. Proos J, Zalka C. Shor's discrete logarithm quantum algorithm for elliptic curves. Quantum Infor- mation & Computation. 2003;3(4):317-344.
56. Groth J. On the size of pairing-based non-interactive arguments. In: EUROCRYPT 2016. Springer; 2016. p. 305-326.
57. Brumley BB, Tuveri N. Remote timing attacks are still practical. In: European Symposium on Research in Computer Security. Springer; 2011. p. 355-371.
58. Biehl I, Meyer B, Müller V. Differential fault attacks on elliptic curve cryptosystems. In: Advances in Cryptology—CRYPTO 2000. Springer; 2000. p. 131-146.
59. National Institute of Standards and Technology. Digital Signature Standard (DSS). Federal Infor- mation Processing Standards Publication 186-5. February 2023.
60. National Institute of Standards and Technology. Recommendations for Discrete Logarithm-based Cryptography: Elliptic Curve Domain Parameters. SP 800-186. February 2023.
61. Webber M, Elfving V, Weidt S, Hensinger WK. The impact of hardware specifications on reaching quantum advantage in the fault tolerant regime. AVS Quantum Science. 2022;4(1):013801.
62. Panteleev P, Kalachev G. Asymptotically good quantum and locally testable classical LDPC codes. In: Proceedings of the 54th Annual ACM Symposium on Theory of Computing. 2022. p. 375-388.
63. Dallaire-Demers PL, Doyle W, Foo T. Brace for Impact: ECDLP Challenges for Quantum Crypt- analysis.

 doi:10.20944/preprints202509.2079.v1

arXiv:2508.14011. August 2025.

64. Fowler AG, Mariantoni M, Martinis JM, Cleland AN. Surface codes: Towards practical large-scale quantum computation. Physical Review A. 2012;86(3):032324.
65. Litinski D. Magic state distillation: Not as costly as you think. Quantum. 2019;3:205.
66. Gidney C, E k e rå M. How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. Quantum. 2021;5:433.
67. Gheorghiu V, Mosca M. Benchmarking quantum cryptanalysis against RSA and ECC. Quantum Science and Technology. 2024;9(2):025003.
68. Bravyi S, Cross AW, Gambetta JM, Maslov D, Rall P, Yoder TJ. High-threshold and low-overhead fault-tolerant quantum memory. Nature. 2023;614:676-681.
69. Clark J. SegWit Transaction Capacity Analysis. Coin Metrics State of the Network. Issue 82. De- cember 2020. Available at: https://coinmetrics.io/segwit-adoption/
70. Bashir A, Paquin C. The Great Bitcoin Scaling Debate: A Technical Post-Mortem. IEEE Interna- tional Conference on Blockchain. 2018. p. 1652-1659.
71. Campbell R. Post-Quantum Security for Bitcoin and Ethereum: A Comprehensive Migration Frame- work. Preprints.org. Posted: August 22 2025. doi:10.20944/preprints202508.1672.v1 [Not peer-reviewed]
72. Bitnodes. Global Bitcoin Nodes Distribution. Retrieved September 20, 2025. Available at: https://bitnodes.io
73. Ethernodes. Ethereum Mainnet Node Statistics. Retrieved September 20, 2025. Available at: https://ethernodes.org
74. Moore S, Anderson R. Economic Barriers to Blockchain Node Operation Under Post-Quantum Cryp- tography. Journal of Cryptoeconomics. 2024;8(2):112-128.
75. Mosca M, Piani M. Quantum Threat Timeline Report 2024. Global Risk Institute. 2024. [Provides attack time calculations for various quantum computer configurations]
76. Espitau T, Fouque PA, G'erard B, Tibouchi M. Side-channel attacks on BLISS lattice-based signa- tures: Exploiting branch tracing against strongSwan and electromagnetic emanations in microcontrollers. ACM Conference on Computer and Communications Security. 2017. p. 1857-1874.