PEER-REVIEWED JOURNAL OF COMPUTER SCIENCE (PRJCS)

ISSN Online : Open Access Scholarly Journal

Post-Quantum Cryptography: Standardisation, Algorithms, and Deployment Migration

Bini P B

Assistant Professor, Department of Computer Science, CCSIT Dr John Matthai Center, Thrissur, India.

Article information

Received: 12th April 2026
Received in revised form: 22nd April 2026
Accepted: 1st May 2026
Available online: 10th May 2026

Volume: 1
Issue: 5
DOI: https://doi.org/10.5281/zenodo.20136246

Abstract

The prospect of cryptographically relevant quantum computers threatens public-key primitives underpinning modern digital infrastructure, including RSA, Diffie-Hellman, and elliptic-curve cryptography. Post-quantum cryptography (PQC) responds with algorithms believed to be hard for both classical and quantum adversaries. In August 2024, the U.S. National Institute of Standards and Technology (NIST) finalised the first three PQC standards: FIPS 203 (ML-KEM, based on Kyber), FIPS 204 (ML-DSA, based on Dilithium), and FIPS 205 (SLH-DSA, based on SPHINCS+). This paper surveys the algorithmic families behind these standards, the harvest-now-decrypt-later threat model, hybrid migration strategies, performance trade-offs, and the engineering challenges of transitioning the global cryptographic stack. We examine deployment progress in TLS, SSH, X.509, and constrained devices, and outline open research questions on side-channel security, falcon-style signatures, and the long-term replacement schedule for classical asymmetric cryptography.

Keywords: Post-quantum Cryptography, Lattice Cryptography, ML-KEM, ML-DSA, SPHINCS+, NIST Standardisation, Harvest Now Decrypt Later, Hybrid TLS.

.

I. INTRODUCTION

Modern asymmetric cryptography rests on the conjectured hardness of integer factorisation and the discrete-logarithm problem in finite fields and elliptic-curve groups. Shor's quantum algorithm [1] solves both problems in polynomial time on a sufficiently large fault-tolerant quantum computer. Although such a machine does not yet exist, the harvest-now-decrypt-later attack in which an adversary records ciphertexts today and decrypts them when quantum hardware matures already threatens long-lived secrets such as state communications, intellectual property, and medical records. Post-quantum cryptography (PQC) provides public-key primitives whose security is believed to hold against both classical and quantum adversaries [2].

After a multi-year competitive process initiated in 2016, NIST finalised the first three PQC standards in August 2024: FIPS 203 specifying ML-KEM (the Kyber-derived key-encapsulation mechanism), FIPS 204 specifying ML-DSA (the Dilithium-derived digital signature algorithm), and FIPS 205 specifying SLH-DSA (the SPHINCS+-derived stateless hash-based signature) [3], [4], [5]. Falcon, a more compact lattice signature, is in the process of being standardised as FIPS 206. The remainder of this paper surveys the underlying algorithm families, performance characteristics, and deployment trajectory.

II. THE QUANTUM THREAT MODEL

Two quantum algorithms are central to the threat model. Shor's algorithm [1] factors integers and computes discrete logarithms in time polynomial in the input size, thereby breaking RSA, DH, and ECC once a fault-tolerant quantum computer of sufficient logical-qubit count exists. Grover's algorithm [6] provides a quadratic speed-up for unstructured search, halving the effective security of symmetric primitives. The practical defence against Grover is

straightforward: doubling the key length restores the original security level. The defence against Shor is more disruptive because no parameter adjustment of RSA or ECC suffices; an algorithmic replacement is required. Resource estimates by Gidney and Ekerå [7] suggest that breaking RSA-2048 with Shor would require on the order of 20 million noisy qubits, well beyond current hardware, but the long lead time of cryptographic migration makes early action prudent.

III. ALGORITHMIC FAMILIES

Four mathematical foundations dominate modern PQC. Lattice-based cryptography exploits the hardness of problems such as Learning with Errors (LWE) [8] and Module-LWE [9]; Kyber [10] and Dilithium [11] are concrete instantiations now standardised as ML-KEM and ML-DSA. Code-based cryptography, originating with McEliece's 1978 cryptosystem [12], builds on the hardness of decoding general linear codes; the Classic McEliece submission remains a candidate for long-term diversity. Hash-based signatures, including stateful schemes such as XMSS [13] and the stateless SPHINCS+ [14], rely only on the security of cryptographic hash functions and provide the most conservative security argument. Multivariate-quadratic and isogeny-based schemes provide additional diversity, although the SIDH/SIKE family was broken in 2022 by Castryck and Decru [15], a striking reminder that PQC security is an active research frontier.

Among the standardised primitives, Falcon is notable for combining NTRU lattices with hashing-and-sampling techniques, yielding small signatures at the cost of complex constant-time floating-point implementation. The selection of multiple signature families (ML-DSA, SLH-DSA, Falcon) reflects NIST's policy of cryptographic diversity: should a flaw be found in one family, deployers can migrate to another without abandoning PQC entirely.

Fig. 1. Public-key and signature/ciphertext sizes (log scale) for classical and standardised PQC primitives.

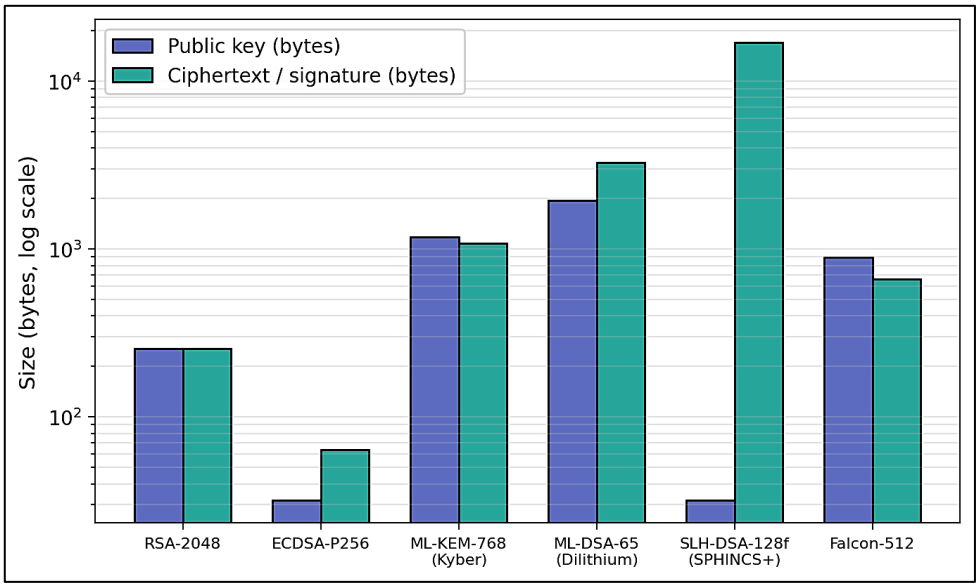

IV. PERFORMANCE AND IMPLEMENTATION TRADE-OFFS

PQC primitives differ from RSA and ECC in characteristic ways. ML-KEM-768 produces a 1184-byte public key and 1088-byte ciphertext, compared with the 256-byte modulus of RSA-2048 and the 32-byte compressed point of ECDH-P256. ML-DSA-65 signatures are roughly 3.3 KB; SPHINCS+ at the small parameter set is closer to 8 KB and reaches 50 KB at higher security levels. These size increases have direct consequences for handshake bandwidth and on-disk certificate storage. Computationally, lattice operations are fast typically a few hundred microseconds and often outperform RSA, but stateless hash-based signatures impose a significant computational cost in exchange for their conservative security argument [16].

Side-channel security is a recurring concern. Constant-time implementations of lattice algorithms are well understood [17], but recent attacks demonstrated practical recovery of secrets from non-constant-time implementations of Kyber and Dilithium. Hardware-accelerated implementations are now appearing in libraries such as liboqs and in mainstream TLS stacks including OpenSSL 3.5 and BoringSSL [18].

V. STANDARDISATION AND DEPLOYMENT

Standardisation has progressed in parallel with experimental deployment. The IETF has finalised the Hybrid Public Key Encryption (HPKE) framework [19] and a series of TLS specifications combining ML-KEM with classical key exchange. Cloudflare and Google reported successful field experiments of hybrid X25519+Kyber TLS handshakes covering tens of percent of HTTPS traffic by mid-2024 [20]. Apple introduced PQ3, a post-quantum key-agreement protocol for iMessage, in 2024 [21]. Major operating systems and browsers ship PQC support in TLS 1.3 by default in 2025-26 releases. The CNSA 2.0 suite of the U.S. National Security Agency mandates PQC for national-security systems by 2035 [22].

Fig. 2. Projected migration share of classical vs. post-quantum public-key cryptography, 2024-2035.

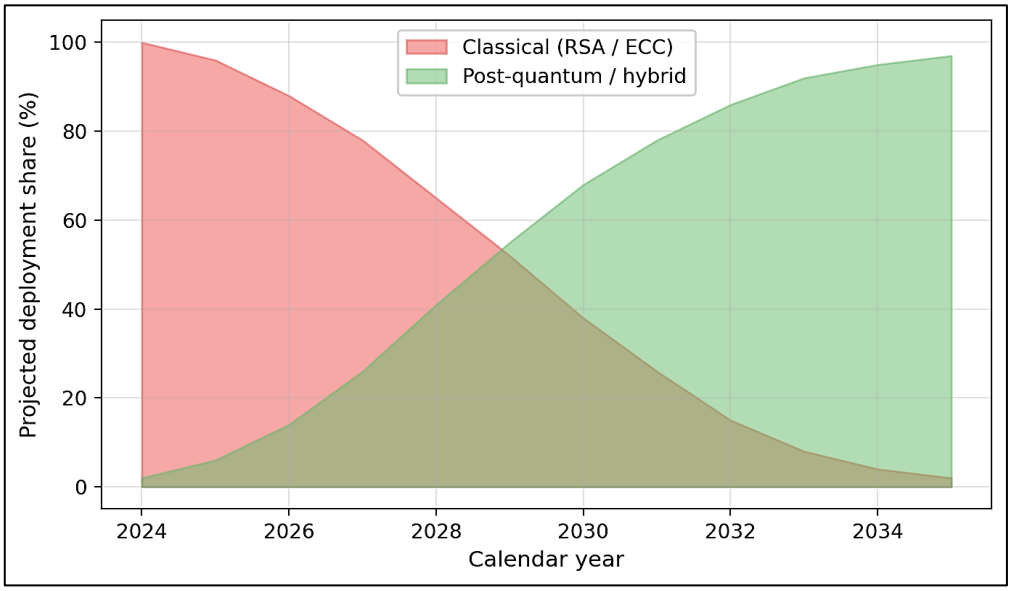

Table 1. NIST Post-Quantum Standards Compared With Classical Public-Key Cryptography

Standard	Type	Underlying problem	Public key (bytes)
FIPS 203 (ML-KEM-768)	KEM	Module-LWE	1184
FIPS 204 (ML-DSA-65)	Signature	Module-LWE / SIS	1952
FIPS 205 (SLH-DSA-128f)	Signature	Hash collision resistance	32
FIPS 206 (Falcon-512, in progress)	Signature	NTRU lattice	897
RSA-2048 (classical)	KEM/Sig	Integer factorisation	256
ECDSA P-256 (classical)	Signature	Discrete log	32

VI. MIGRATION STRATEGIES

Cryptographic migrations historically take a decade or more to complete; the move from MD5 to SHA-2 and from SHA-1 to SHA-256 illustrates the inertia of deployed systems. PQC migration is harder because public keys, certificates, hardware tokens, smart cards, and embedded devices must all be replaced. Three migration strategies dominate. Hybrid deployment combines a classical and a post-quantum primitive in parallel, providing security as long as either holds; this is the dominant strategy for TLS in 2024-26. Pure-PQC deployment replaces classical primitives outright and minimises bandwidth, but discards classical fallback if a flaw is found in the chosen PQC scheme. Crypto-agility engineering systems so that primitives can be swapped without protocol changes is the more general goal, since further algorithmic replacements are likely as the field matures [23], [24].

VII. OPEN PROBLEMS

Several research questions remain. The lattice problems underpinning ML-KEM and ML-DSA have withstood decades of cryptanalysis but lack the long history of integer factorisation; conservative users may prefer hash-based or code-based signatures despite their size penalty. Side-channel resistance of constant-time PQC on constrained devices, including smart cards and IoT chips, is still an active research area. Authenticated key exchange protocols beyond TLS including DNSSEC, code signing, and email face their own migration constraints. Finally, the governance of the next round of NIST candidates, the role of code-based primitives such as HQC and BIKE, and the long-term coexistence of PQC with classical cryptography all remain open [25].

VIII. ENGINEERING THE TRANSITION IN PRACTICE

The first practical step in any post-quantum migration is an inventory. Organisations rarely have a complete record of where classical asymmetric cryptography is used inside their systems, and this is the single largest source of avoidable cost in PQC programmes. TLS endpoints, code signing, document signing, hardware tokens, smart cards, VPN gateways, certificate authorities, S/MIME and PGP implementations, and embedded firmware update channels all need to be enumerated. Tools such as the Open Quantum Safe project's discovery utilities, Microsoft's PQC inventory libraries, and bespoke scanners commissioned by financial-sector regulators have begun to address this gap. Without a credible inventory, schedule estimates for the migration are guesses, and the schedule cost of guessing wrong is significant.

The second practical issue is hardware. Many production deployments rely on hardware security modules and smart cards whose firmware does not yet support PQC algorithms. The lattice schemes are computable on contemporary HSMs, but ML-DSA signatures at three kilobytes per signature exceed the storage budgets of legacy PIV smart-card slots and many embedded TPMs. SLH-DSA at the higher security parameter sets reaches sizes that make smart-card use impractical without redesigning the surrounding protocol. Vendors are responding with new hardware lines, but

procurement and deployment cycles for hardware run on five- to ten-year horizons, comparable to the cryptographically relevant quantum-computer timeline. The migration cannot be deferred indefinitely without accepting hardware-refresh risk.

Hybrid deployment in TLS has progressed faster than other layers because TLS terminates at software endpoints with controlled refresh cycles. Cloudflare reported that, by the second half of 2024, a substantial fraction of HTTPS traffic from supporting browsers used the X25519+Kyber hybrid key exchange [20], and the share has continued to climb. Apple's PQ3 protocol for iMessage [21] illustrates a different deployment model in which an end-to-end application provides its own post-quantum guarantees independent of the underlying transport. Both deployments are encouraging proofs of concept, but they leave the bulk of the cryptographic surface area, including code-signing chains, X.509 root programmes, and DNSSEC, almost untouched. The next two to three years will likely see those gaps close in turn.

The lasting investment, beyond the migration itself, is crypto-agility. Systems designed to negotiate primitives, retire deprecated ones, and roll forward to replacements without protocol redesign will weather the inevitable second migration when, in the future, a flaw is found in one of the present standards or a more efficient alternative becomes available. The IETF Hybrid Public Key Encryption framework [19], TLS 1.3 cipher-suite negotiation, and modern certificate transparency logs each contribute to this property. Greenfield systems built today should treat algorithm choice as a parameter, not a hard-coded constant. Legacy systems are harder to retrofit, but the cost of crypto-agility retrofit is, on the whole, lower than the cost of repeated emergency migrations driven by the next round of cryptanalytic surprises.

A perspective worth keeping in mind is that the cryptographically relevant quantum computer remains a moving target, and serious estimates of when it might arrive vary by decades depending on whose roadmap one trusts. The conservative engineering response is not to predict the arrival date precisely but to ensure that long-lived secrets are protected against a plausible date range that includes the next ten to twenty years. From that vantage point, the specific timing of when Shor becomes feasible is less important than whether one's certificates, signed firmware, archived medical records, and signed legal documents will outlast the threat. Several governments and standards bodies have converged on a target window of 2030-2035 for substantial completion of the migration, and that window is treated as a regulatory expectation rather than a guess. The PQC standards finalised in 2024 give organisations the algorithms they need; what is now lacking is execution capacity at the scale required, including trained engineers, vendor-supported tooling, and procurement budgets aligned to the timeline. The next few years will reveal whether organisational readiness can match algorithmic readiness, and recent surveys suggest that many sectors are still at the inventory stage. The risk for laggards is not that they fail to adopt PQC at all, but that they will be forced to adopt it under emergency conditions when news of a successful attack against classical primitives breaks. The cost of a planned migration is predictable; the cost of an unplanned one rarely is.

IX. CONCLUSION

The completion of FIPS 203, 204, and 205 in 2024 marks the start, not the end, of the post-quantum transition. PQC primitives are now sufficiently mature for production deployment, particularly in hybrid configurations alongside classical cryptography. The substantial work that lies ahead migrating certificate authorities, hardware security modules, embedded devices, and long-lived protocols will dominate cryptographic engineering for the next decade. Crypto-agility, conservative algorithm choice, and disciplined incident response will determine whether the global digital infrastructure can be safely transitioned before cryptographically relevant quantum computers arrive.

REFERENCES

[1] P. W. Shor, “Algorithms for quantum computation: Discrete logarithms and factoring,” in *Proc. IEEE Symp. Foundations of Computer Science (FOCS)*, 1994.

[2] D. J. Bernstein and T. Lange, “Post-quantum cryptography,” *Nature*, vol. 549, pp. 188–194, 2017.

[3] National Institute of Standards and Technology, *FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard*. Gaithersburg, MD, USA: NIST, Aug. 2024.

[4] National Institute of Standards and Technology, *FIPS 204: Module-Lattice-Based Digital Signature Standard*. Gaithersburg, MD, USA: NIST, Aug. 2024.

[5] National Institute of Standards and Technology, *FIPS 205: Stateless Hash-Based Digital Signature Standard*. Gaithersburg, MD, USA: NIST, Aug. 2024.

[6] L. K. Grover, “A fast quantum mechanical algorithm for database search,” in *Proc. ACM Symp. Theory of Computing (STOC)*, 1996.

[7] C. Gidney and M. Ekerå, “How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits,” *Quantum*, vol. 5, p. 433, 2021.

[8] O. Regev, “On lattices, learning with errors, random linear codes, and cryptography,” *J. ACM*, vol. 56, no. 6, 2009.

[9] A. Langlois and D. Stehlé, “Worst-case to average-case reductions for module lattices,” *Des. Codes Cryptogr.*, vol. 75, pp. 565–599, 2015.

[10] J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck *et al.*, “CRYSTALS-Kyber: A CCA-secure module-lattice-based KEM,” in *Proc. IEEE Eur. Symp. Security Privacy (EuroS&P)*, 2018.

[11] L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé, “CRYSTALS-Dilithium: A lattice-based digital signature scheme,” *IACR Trans. Cryptogr. Hardw. Embedded Syst.*, 2018.

[12] R. J. McEliece, “A public-key cryptosystem based on algebraic coding theory,” *Deep Space Netw. Prog. Rep.*, vol. 44, pp. 114–116, 1978.

[13] J. Buchmann, E. Dahmen, and A. Hülsing, "XMSS—A practical forward secure signature scheme based on minimal security assumptions," in *Proc. PQCrypto*, 2011.
[14] D. J. Bernstein, A. Hülsing, S. Kölbl, R. Niederhagen, J. Rijneveld, and P. Schwabe, "The SPHINCS+ signature framework," in *Proc. ACM Conf. Comput. Commun. Security (CCS)*, 2019.
[15] W. Castryck and T. Decru, "An efficient key recovery attack on SIDH," in *Proc. EUROCRYPT*, 2023.
[16] G. Alagic, J. Alperin-Sheriff, D. Apon, D. Cooper, Q. Dang, J. Kelsey *et al.*, *Status Report on the Third Round of the NIST Post-Quantum Cryptography Standardization Process*, NISTIR 8413, 2022.
[17] P. Pessl, L. G. Bruinderink, and Y. Yarom, "To BLISS-B or not to be: Attacking strongSwan's implementation of post-quantum signatures," in *Proc. ACM Conf. Comput. Commun. Security (CCS)*, 2017.
[18] Open Quantum Safe Project, "liboqs: C library for prototyping and experimenting with quantum-resistant cryptography," GitHub repository, 2024.
[19] R. L. Barnes, K. Bhargavan, B. Lipp, and C. A. Wood, *Hybrid Public Key Encryption*, RFC 9180, IETF, 2022.
[20] K. Kwiatkowski, N. Sullivan, A. Langley, D. Levin, and A. Mislove, "Measuring TLS key exchange with post-quantum KEM," in *Proc. NIST PQC Workshop*, 2019.
[21] Apple Security Engineering and Architecture, "iMessage with PQ3: The new state of the art in quantum-secure messaging at scale," *Apple Security Research Blog*, Feb. 2024.
[22] U.S. National Security Agency, *Commercial National Security Algorithm Suite 2.0 (CNSA 2.0)*, Cybersecurity Advisory CSA-22-1158, 2022.
[23] L. Chen, S. Jordan, Y.-K. Liu, D. Moody, R. Peralta, R. Perlner, and D. Smith-Tone, *Report on Post-Quantum Cryptography*, NISTIR 8105, 2016.
[24] M. Mosca, "Cybersecurity in an era with quantum computers: Will we be ready?," *IEEE Security Privacy*, vol. 16, no. 5, pp. 38–41, 2018.
[25] D. Joseph, R. Misoczki, M. Manzano, J. Tricot, F. D. Pinuaga, O. Lacombe, S. Leichenauer *et al.*, "Transitioning organizations to post-quantum cryptography," *Nature*, vol. 605, pp. 237–243, 2022.