International Journal of
Technology and Emerging Research
ISSN: 3068-109X (online) | ISSN: 3068-1995 (print) | https://www.ioro.org/ijter

International Conference on Intelligent Technologies for Tomorrow (ICITT-2026) · 18 August 2026

Post-Quantum Cryptographic Standards and Migration Obligations: A Review of Status, Cryptanalysis, and Policy

Sajeevkrishnan M. A. [1], **Arifa P.** [1,*]

[1]*Department of Computer Science and Application, CHMKM Government Arts and Science College, Tanur, Kerala, India*

* Corresponding author: sajeevkrishnan@gmail.com, arifanoushadcn@gmail.com

Abstract:

Quantum computing threatens the public-key cryptography underpinning digital infrastructure, and the practical question has shifted from whether that threat will materialise to how quickly organisations must respond. Two developments drove the shift: the first post-quantum cryptographic standards were published in August 2024, and by 2026 several governments had attached dated deadlines to their adoption. This review synthesises primary standards and policy documents, measured deployment data, and recent research literature. Four findings emerge. First, the standardised portfolio is smaller than commonly reported: three standards are published, a fourth remains in development, and a fifth algorithm has been selected but not standardised, a distinction that matters because procurement obligations attach only to published standards. Second, migration mandates are now international and broadly convergent, with the United States, United Kingdom and Canada all terminating in 2035, while India targets full adoption for critical infrastructure by 2029; these instruments differ in legal force, so their dates are not equivalent obligations. Third, measured deployment shows the transition far advanced where few parties control both endpoints and barely begun in the interior, with more than half of client-facing web traffic protected but only a small fraction of origin servers. Fourth, two candidate algorithms that reached advanced evaluation were defeated by classical rather than quantum cryptanalysis, one within a weekend on consumer hardware. The review concludes that cryptographic agility, not any particular algorithm, is the durable objective, and that migration should be prioritised by the confidentiality lifetime of data rather than by forecasts of adversary capability.

Keywords: post-quantum cryptography; cryptographic migration; cryptographic agility; standardization; critical infrastructure

1. Introduction

The question facing organisations that depend on public-key cryptography is no longer whether quantum computing will invalidate it, nor even when, but on what schedule they are obliged to respond. Between August 2024 and June 2026 that question acquired concrete answers from two directions at once. The United States published its first post-quantum cryptographic standards, and it then attached binding dates to their adoption.

The technical milestone came first. Three Federal Information Processing Standards — FIPS 203, FIPS 204, and FIPS 205 — were published on 13 August 2024, specifying algorithms derived from CRYSTALS-Kyber, CRYSTALS-Dilithium, and SPHINCS+ respectively [1]. A fourth algorithm, Falcon, was selected but "will be published in FIPS 206 (in development)" and remains unpublished [1], while HQC was selected for standardization on 11 March 2025 [1] and has yet to receive a FIPS number. The available toolkit is therefore

smaller than the frequently repeated count of standardised algorithms implies, a discrepancy this review treats as consequential rather than pedantic.

The policy milestone followed. Executive Order 14412, signed 22 June 2026, requires federal agencies to transition all high value assets and high impact systems to post-quantum key establishment by 31 December 2030 and to post-quantum digital signatures by 31 December 2031, and directs the Federal Acquisition Regulatory Council to propose a rule obliging covered contractors to comply with post-quantum FIPS by the earlier of those dates [2]. Implementing guidance sets out a five-phase migration schedule running to 2035 [3]. Migration has thus moved from advisory to mandatory, and its reach extends past government to any organisation that sells to it. Nor is this a single-jurisdiction phenomenon. By 2026 national migration guidance had also been issued in the United Kingdom [19], Canada [20], India [17], and the European Union [21], on timetables that broadly converge on 2035 but through instruments that differ substantially in legal force — and, in India's case, on a markedly faster schedule for critical infrastructure.

Underlying both developments is a threat that does not wait for the hardware. As the order itself states, adversary activity "presents the risk of adversaries collecting United States information now, and decrypting it later once large-scale quantum computers are operational" [2]. Any data whose confidentiality must outlive the arrival of a cryptographically relevant quantum computer is already exposed, which is why federal guidance prioritises systems containing "data expected to remain mission-sensitive in 2030" [3] rather than systems judged fragile today.

This review proceeds as follows. It first sets out how its sources were gathered, then surveys the research literature published since standardization, which has shifted its subject from algorithm design to deployment engineering. It then establishes the background and the state of quantum hardware and resource estimation, and examines the threat model that makes the migration timetable urgent. The central sections address what has actually been standardised and what has not, how the resulting algorithms compare against each other and against what they replace, what obligations now bind and whom they reach across five jurisdictions, and what measured deployment in production systems reveals about feasibility. A discussion of transition challenges follows, then specialized applications, open research problems, and conclusions.

Its emphasis throughout is on distinctions that determine what an organisation can lawfully procure and deploy today: between algorithms that are published, selected, or merely announced; between mandates that bind and guidance that recommends; and between deployment that has been measured and adoption that has been announced.

2. Review Methodology

This review combines a primary-source study of standards and policy instruments with a search of the research literature. Standards and policy documents were retrieved directly from their issuing bodies: the NIST Computer Security Resource Center and the NIST public-access repository for Federal Information Processing Standards and Internal Reports; the *Federal Register* for the text of executive orders as published; and the Executive Office of the President for implementing memoranda. National migration guidance was retrieved from each issuing authority rather than through comparative secondary summaries — the Department of Science and Technology for India, the National Cyber Security Centre for the United Kingdom, the Canadian Centre for Cyber Security, and the *Official Journal* for European Union instruments. Preprints and conference papers were retrieved from arXiv and the IACR Cryptology ePrint Archive. Bibliographic details — venue, pagination, and digital object identifier — were confirmed against Crossref records.

The research literature was searched on 9 August 2026 against the arXiv e-print archive, using its programmatic interface with the query terms "post-quantum" and "migration" restricted to abstracts and ordered by submission date, and against the IACR Cryptology ePrint Archive using the same terms. Searches of IEEE Xplore, the ACM Digital Library, Scopus, and DBLP were attempted and did not complete: the first three refused automated queries and the fourth reported its search service unavailable. Literature indexed exclusively in those repositories was therefore not searched systematically. Peer-reviewed work cited here was instead identified through the two accessible archives and through general search, and each item was then verified against its publisher record. This is a weaker instrument than a full database sweep, since it preferentially surfaces work that is also available as a preprint, and the limitation is stated again below.

The principal date range is January 2022 to August 2026, spanning the final NIST selection rounds, publication of the first standards, and the emergence of binding migration mandates. Foundational works are included without date restriction where a primary citation is required.

Four source classes were admitted: peer-reviewed publications; official standards and government policy documents; preprints of resource-estimation work, cited as preprints where no peer-reviewed version exists; and operator engineering disclosures reporting measured data from systems the disclosing party itself runs. The last class is admitted on narrow terms. A disclosure qualifies only where it reports quantities the operator has measured — traffic shares, handshake latencies, failure rates — and it is cited for those measurements alone and not for the surrounding commentary. Vendor announcements, product availability and readiness claims, roadmap statements, and unquantified assertions of support are excluded whatever their origin, as are press releases, undated web content, and secondary sources restating primary claims. The exclusion of secondary sources was applied even where the restating document was itself official: several national roadmaps survey other jurisdictions' timelines, and those surveys were used to locate primary sources rather than cited in their place.

Three sourcing constraints were applied. Standards and policy documents were read in primary form rather than through secondary coverage, a practice that proved necessary — reporting on the June 2026 executive order disagreed on its publication date, which the *Federal Register* record resolved. Quantitative claims, including algorithm parameter sizes, attack costs, and quantum resource estimates, were taken only from the documents stating them, and derived figures are identified as such. Where a policy document advances a threat-horizon estimate sourced to executive commentary rather than technical analysis, that estimate is reported as the document's stated rationale and is not adopted as a finding of this review.

Four limitations follow. Coverage of the peer-reviewed literature is incomplete for the reason given above: four bibliographic databases could not be queried, so the search rests on two open preprint archives supplemented by individual verification, and work that is peer-reviewed but never preprinted is liable to be under-represented. Deployment evidence is drawn partly from operator disclosures, which exist only where an operator chooses to publish and which characterise that operator's vantage point rather than the Internet as a whole; an independent multi-vantage measurement study is used to check the operator figures where the two overlap, but both sources observe web traffic and neither reaches enterprise interior networks, operational technology, or protocols other than the web. Enterprise migration cost data remains largely proprietary; a published governmental estimate exists for one jurisdiction's civilian systems and is reported in the discussion of cost, but no comparable cross-industry figure could be sourced, and cost drivers are therefore reported in place of estimates. Coverage of national migration guidance is confined to jurisdictions publishing in English or in English translation, and to those whose instruments were retrievable in primary form; several further national programmes were identified but not obtained, and are not represented.

3. Related Literature: From Algorithm Design to Deployment Engineering

The research literature on post-quantum cryptography has undergone a change of subject since the 2018 review this paper updates [31], and characterising that change is the clearest way to situate the present contribution. That earlier survey set out to "elucidate the implications of quantum computing in present cryptography and to introduce the reader to basic post-quantum algorithms" [31], surveying lattice-, multivariate-, hash-, and code-based families as candidates. Each of those families has since been resolved by the standardization process described below: lattice and hash constructions were standardised, a code-based scheme was selected, and the multivariate finalist was broken outright. The questions that remain are no longer about which families are viable. Work published before standardization was largely concerned with algorithm design and cryptanalysis — with whether candidate schemes were secure and how fast they could be made. The work surfaced by the searches described in the methodology is instead concerned with deployment: with measuring adoption, securing implementations, verifying protocols that compose the new primitives, and building the tooling migration requires. The threat is no longer the open question. The engineering is.

Four strands are active, and each bears directly on a section of this review.

The first is **empirical measurement of deployment**, which barely existed before the standards did. The most substantial contribution is a longitudinal study establishing "more than 2 billion TLS handshakes" and analysing "cryptographic negotiation behavior across 1 million domains from 11 globally distributed vantage points" [29]. Its findings are consequential for any policy-centred account of the migration, including this one. It reports "configuration convergence: PQ-TLS deployment overwhelmingly centers on a single hybrid construction," and that "much of the apparent progress is driven by managed infrastructure providers" — an independent confirmation, at far greater scale, of the edge-versus-interior asymmetry discussed in the deployment section below. More pointedly, it finds that "[n]ational timelines and sectoral priorities show limited correspondence with observed deployment patterns" [29]. That is a measured challenge to the assumption that mandates drive adoption, and it is engaged rather than deflected in the discussion of migration mandates.

The second strand is **implementation security**, which has become the live attack surface now that the algorithms are settled. Current work targets protected implementations rather than naive ones: recent analysis of "the side-channel vulnerability of masked ML-DSA implementations" demonstrates first-, second-, and higher-order attacks against the masking countermeasure itself [24]. The significance is the shift in what is being attacked. The pre-standardization literature broke schemes; the post-standardization literature breaks constructions of schemes that remain mathematically sound.

The third strand is **formal verification of protocols** that compose post-quantum primitives. Verification of Signal's PQXDH protocol identified "several flaws and potential vulnerabilities in the PQXDH specification," none exploitable in the deployed application, and established that proving the deployment secure required "an additional binding property of the KEM" that no standard supplied [22]. This literature is small but disproportionately informative, because it examines the joint between a standardized primitive and the protocol that carries it — which is where the standards stop and deployment begins.

The fourth strand is **migration and measurement methodology**, including the unglamorous problem of comparing implementations honestly. Recent work on ML-DSA benchmarking argues that the algorithm's "execution-time variability … an inherent property due to rejection sampling" makes conventional single-figure metrics unsuitable for engineering decisions, and proposes standardized data sets and qualified

reporting instead [30]. This review adopts that caution directly: it declines to construct a cross-algorithm performance table from measurements taken on different platforms by different parties.

Two observations follow, with the caveat that they characterise the work this review was able to reach rather than the field exhaustively. The literature examined has become measurement-led, which is a mark of maturity — claims about adoption, overhead, and failure are increasingly settled by observation rather than by extrapolation from microbenchmarks. And it has begun to disagree productively with policy: the measurement studies find weaker correspondence between mandates and deployment than the mandates themselves assume. A review that reported only the standards and the timetables would miss this, which is one reason the present paper treats measured deployment as a distinct body of evidence rather than as illustration.

The coverage of this section is subject to the search limitation recorded in the methodology. Four bibliographic databases could not be queried, so the literature reviewed here was reached through open preprint archives and individual verification. Work that is peer-reviewed but never preprinted is liable to be under-represented, and this section should be read as a characterisation of active research directions rather than as an exhaustive survey.

4. Background: From Candidate Algorithms to Binding Deadlines

The cryptography now scheduled for replacement is not confined to specialist systems. Federal guidance describes the dependency in ordinary terms: "[f]or decades, strong cryptography has enabled the United States Government to protect Federal information, securely deliver critical services to the American people, and guard against cyber-enabled fraud," and notes that citizens "depend heavily on encryption to protect their individual privacy and for everyday tasks such as starting their cars, paying for groceries, and messaging friends and family" [3]. This review assumes familiarity with RSA and elliptic-curve cryptography and with the hardness assumptions on which they rest; what matters for the argument is that these primitives are embedded in payment infrastructure, vehicle immobilisers, messaging, and identity systems, and that replacement is therefore an estate-wide engineering exercise rather than a cryptographic one. The specific algorithms placed at risk, and the mechanism by which they fail, are treated in Section 6.

Replacement candidates have been under evaluation for a decade. NIST initiated the Post-Quantum Cryptography Standardization Process in December 2016, and after three rounds announced its first selections: CRYSTALS-Kyber as a key-encapsulation mechanism, with CRYSTALS-Dilithium, Falcon, and SPHINCS+ as digital signature schemes [4]. Four additional key-encapsulation candidates — BIKE, Classic McEliece, HQC, and SIKE — entered a fourth round in July 2022, selected because they "were all based on different security assumptions than ML-KEM" [4]. Implementing guidance characterises the exercise as one in which "many of the world's most respected cryptographers and security researchers have thoroughly evaluated candidate algorithms to determine whether they could be compromised or broken" by a cryptographically relevant quantum computer [3].

Two developments converted that evaluation into an obligation. FIPS 203, 204, and 205 were published on 13 August 2024 [1], and the fourth round closed with the selection of HQC on 11 March 2025 [1]. Then, on 22 June 2026, Executive Order 14412 attached dates to adoption, requiring post-quantum key establishment across high value assets and high impact systems by 31 December 2030 and post-quantum digital signatures by 31 December 2031 [2]. The interval between a published standard and a mandated deadline was under two years. That compression, rather than any change in the underlying mathematics, is what distinguishes

the current moment: the question has shifted from which algorithms will be available to how quickly existing estates can be made to use them.

5. Error Correction, Resource Estimates, and Migration Timelines

What the Published Evidence Supports

This review does not survey quantum hardware announcements, whose competing claims are difficult to evaluate against a consistent standard and which fall outside the source classes set out above. It relies instead on peer-reviewed error-correction results, published resource estimates, and the assessment of the body mandating the migration. Vendor development roadmaps, which state targets rather than results, are excluded on the same basis.

The peer-reviewed result that bears most directly on feasibility concerns error correction rather than qubit count. Google Quantum AI reported two surface-code memories operating below the error-correction threshold, with the logical error rate of the larger memory "suppressed by a factor of $\Lambda = 2.14 \pm 0.02$ when increasing the code distance by two" [32]. The distance-7 memory used 101 qubits and achieved an error rate of $0.143\% \pm 0.003\%$ per cycle, exceeding the lifetime of its best constituent physical qubit "by a factor of 2.4 ± 0.3" [32]. The significance for migration planning is qualitative rather than predictive: below-threshold operation is the property that makes adding physical qubits reduce logical error rather than increase it, and it is therefore the precondition for the scaling that the resource estimates below assume. Demonstrating it at distance seven is a long way from the scale those estimates require, but it removes a question mark over whether the scaling path exists at all.

Against that, the assessment of the body mandating the migration remains the operative planning statement. Implementing guidance describes a cryptographically relevant quantum computer as a machine that "will be able to decrypt data protected by many forms of cryptography that are commonly used today and thwart existing authentication protocols," and records that "[a] CRQC is not yet known to exist, but steady advancements in the quantum computing field may yield a CRQC in the coming decade" [3]. Two things follow. The capability does not exist as of this writing, on the assessment of the body imposing the deadlines; and the planning horizon adopted by that body is a decade, not a generation.

Resource Estimates for RSA-2048

Resource estimates translate that error-correction picture into a requirement, and they are the most disciplined measure of distance to capability available, because they state their assumptions explicitly and can be compared across time.

Gidney and Ekerå estimated that a 2048-bit RSA integer could be factored in eight hours using twenty million noisy qubits, assuming a planar grid of nearest-neighbour qubits, a physical gate error rate of 10^{-3}, a surface code cycle time of one microsecond, and a reaction time of ten microseconds [5]. The twenty-million figure counts *physical* qubits. The same work states the requirement in the abstract circuit model — which, in the authors' words, "ignores overheads from distillation, routing, and error correction" — as $3n + 0.002n \lg n$ logical qubits [5]. The difference between the two counts is error-correction overhead, and conflating them, as secondary accounts frequently do, overstates the logical requirement by orders of magnitude.

Gidney subsequently estimated that the same factorization could be performed with fewer than one million noisy qubits in under a week [6]. Presenting this as a straightforward reduction in difficulty would be misleading: the qubit requirement fell roughly twentyfold while the runtime rose comparably, making the

revision a space–time tradeoff rather than a uniform improvement. Its significance lies in which resource moved. Qubit count is the harder engineering constraint, so a machine needing a week but fewer than a million qubits sits materially closer to feasibility than one needing twenty million for eight hours.

Planning Under Uncertainty

The lesson of the preceding subsection is not a date. It is that an estimate of this kind moved by more than an order of magnitude in one dimension within six years, on published assumptions and by one of the same authors. Any single figure — including the current one — should be treated as provisional, and migration schedules anchored to a specific predicted arrival are anchored to a moving quantity.

Federal policy resolves this by declining to forecast. Rather than estimating when a CRQC will exist, it prioritises systems by the confidentiality lifetime of the data they hold, directing agencies to attend to those containing "data expected to remain mission-sensitive in 2030" [3]. This substitutes a property of the asset, which an organisation can determine, for a property of an adversary's future capability, which it cannot. It is the most transferable element of the federal approach and is developed further in Section 8.

6. Quantum Threats and Harvest Now, Decrypt Later

Shor's Algorithm and Public-Key Cryptography

The threat to deployed public-key cryptography originates in Shor's polynomial-time quantum algorithms for integer factorization and discrete logarithms [7]. The consequence, rather than the mathematics, is what concerns the enterprise reader: the hardness assumptions underlying RSA and elliptic-curve cryptography do not survive a sufficiently capable quantum computer, and those two families secure the overwhelming majority of key establishment and authentication in production today.

United States federal guidance enumerates the exposure precisely. Elliptic Curve Diffie-Hellman, Menezes-Qu-Vanstone, finite-field Diffie-Hellman, and RSA key establishment are all classified as quantum-vulnerable asymmetric algorithms used for key establishment, while ECDSA, RSA signatures, and DSA are classified as quantum-vulnerable algorithms used for digital signatures [3]. The same guidance instructs agencies to "treat as quantum-vulnerable any asymmetric algorithm that is not definitively known to be quantum-resistant" [3] — a default that inverts the usual burden of proof and is worth adopting in private-sector inventories.

The scale of machine required to realise this exposure, and the extent to which published estimates of it have moved, are treated in Section 5.

Grover's Algorithm and Symmetric Systems

Symmetric cryptography is affected but not broken, and the difference is one of degree rather than kind. Grover's search algorithm [8] provides a quadratic speedup for unstructured search; applied to key recovery, the attack "uses $O(\sqrt{N})$ calls to the cipher to search a key space of size N" [9]. Read naively, this halves the exponent of the search cost and motivates the familiar recommendation to double symmetric key lengths.

That reading treats query count as the cost, which it is not. The realised expense of a Grover attack is governed substantially by the depth of the circuit that must be executed coherently, and it is this constraint rather than the query bound that determines practical feasibility. Jaques et al. "study the cost of quantum key search attacks under a depth restriction and introduce techniques that reduce the oracle depth, even if it requires more qubits" [9], establishing depth as the operative design parameter and trading it against qubit count in the same manner observed for factorization earlier in this section.

The framing matters beyond symmetric ciphers, because it underpins the security categories used throughout the post-quantum standards. As the same authors note, NIST "security categories are defined based on the concrete cost of quantum key search against AES," for which they "present new, lower cost estimates for each category," with "immediate implications for the security assessment of post-quantum cryptography" [9]. The direction of that revision deserves emphasis: the estimates moved lower, not higher. This literature refines the accounting on which the ML-KEM and SLH-DSA category assignments described in Section 7.2 rest; it does not license complacency about symmetric parameters.

Harvest Now, Decrypt Later

The threat that makes the migration timetable urgent does not require a quantum computer to exist yet. Executive Order 14412 states the mechanism in its opening section: "Ongoing cyber activity against our Nation also presents the risk of adversaries collecting United States information now, and decrypting it later once large-scale quantum computers are operational" [2]. Encrypted traffic captured today can be retained and decrypted whenever the capability arrives, which means the exposure window for any given communication opened when it was transmitted, not when a cryptographically relevant quantum computer is finally built.

Federal guidance describes that machine as one which "will be able to decrypt data protected by many forms of cryptography that are commonly used today and thwart existing authentication protocols," and observes that although "[a] CRQC is not yet known to exist … steady advancements in the quantum computing field may yield a CRQC in the coming decade" [3]. The planning consequence follows directly: the relevant question for any dataset is not whether a quantum computer exists during its transmission but whether its confidentiality must outlive the CRQC's arrival.

That reasoning is now embedded in binding federal policy rather than left to risk managers. Agencies are directed to prioritise for migration those systems which "[c]ontain data expected to remain mission-sensitive in 2030" [3] — a criterion defined by the lifetime of the data rather than the age or criticality of the system holding it. Executive Order 14412 attaches dates to the resulting obligation, requiring transition of all high value assets and high impact systems to post-quantum key establishment by 31 December 2030 and to post-quantum digital signatures by 31 December 2031 [2]. Organisations retaining records with multi-decade confidentiality requirements — medical histories, classified material, long-lived financial and legal records — are already inside the exposure window that these deadlines are designed to close.

Public Key Infrastructure and the Authentication Surface

Discussion of the quantum threat tends to centre on confidentiality, but the authentication consequences are at least as severe and arrive on the same schedule. Federal guidance is explicit that after defeating existing protections, "a CRQC could take control of or impersonate devices, systems, and people" [3]. Because ECDSA, RSA signatures, and DSA are all enumerated as quantum-vulnerable [3], the certificate hierarchies, code-signing infrastructure, and device attestation that depend on them are exposed to forgery rather than merely to disclosure — a failure mode that no amount of retrospective re-encryption can repair.

The separation of the two deadlines in Executive Order 14412, with key establishment required a full year ahead of digital signatures [2], reflects the different character of the two problems: harvested ciphertext is a stock of accumulated risk, whereas signature forgery is a flow that begins only when the capability exists. One argues for migrating early; the other permits sequencing.

7. The Standardized Portfolio: A Critical Review

The Standardization Process and Its Conclusion

NIST initiated the Post-Quantum Cryptography Standardization Process in December 2016 to select quantum-resistant public-key algorithms in response to advances in quantum computing [4]. After three rounds of evaluation, NIST announced its first selections: CRYSTALS-Kyber as a key-encapsulation mechanism, and CRYSTALS-Dilithium, Falcon, and SPHINCS+ as digital signature schemes [4]. Four further key-encapsulation candidates — BIKE, Classic McEliece, HQC, and SIKE — advanced to a fourth round that began in July 2022. NIST's stated reason for continuing that round is central to understanding the resulting portfolio: these algorithms "were all based on different security assumptions than ML-KEM" [4]. The fourth round was, in other words, an exercise in insurance rather than optimisation.

That round closed with the selection of HQC in March 2025, which NIST describes as marking "an end to the standardization process which began with the NIST Call for Proposals in 2016" [4]. The qualification NIST attaches to its own conclusion matters for any organisation planning around the results: "not all NIST PQC standardization is concluded, as NIST is also currently evaluating additional digital signatures" [4]. The standard set is complete in the sense that the original process has ended, and incomplete in the sense that two of its outputs have not yet been published as standards.

The Published Standards

Three Federal Information Processing Standards were published on 13 August 2024 [1]. Their status, together with that of the two algorithms still in progress, is summarised in Table 1. It is the single most frequently misstated fact in practitioner literature on this subject, and precision here is not pedantry: procurement obligations attach to published FIPS and not to selections.

Table 1. Status of algorithms from the NIST Post-Quantum Cryptography Standardization Process, as of August 2026.

FIPS	Algorithm	Mathematical basis	Status	Date
FIPS 203	ML-KEM	Module-lattice	Published	13 Aug 2024
FIPS 204	ML-DSA	Module-lattice	Published	13 Aug 2024
FIPS 205	SLH-DSA	Hash-based, stateless	Published	13 Aug 2024
FIPS 206	FN-DSA (Falcon)	NTRU lattice (FFT, hash-and-sign)	In development	Not published
None assigned	HQC	Code-based, quasi-cyclic codes	Selected, not yet standardised	Selected 11 Mar 2025

FIPS 203 specifies ML-KEM, the Module-Lattice-Based Key-Encapsulation Mechanism Standard, in three parameter sets — ML-KEM-512, ML-KEM-768, and ML-KEM-1024 — ordered by "increasing security strength and decreasing performance" [10]. Encapsulation keys occupy 800, 1,184, and 1,568 bytes respectively; decapsulation keys 1,632, 2,400, and 3,168 bytes; ciphertexts 768, 1,088, and 1,568 bytes; and the shared secret key 32 bytes in every case [10]. Two observations follow. First, the standard explicitly declines the conventional shorthand, stating that "security strength is not described by a single number, such as '128 bits of security'" and instead assigning parameter sets to NIST security categories [10]. Papers that render ML-KEM's security levels as bit-equivalents contradict the standard they cite. Second, the sizes are modest in absolute terms but material relative to what they replace; federal implementation guidance characterises ML-KEM as carrying "[l]arger overhead than classical Elliptical Curve Technology (ECC)" [3].

FIPS 204 specifies ML-DSA, the Module-Lattice-Based Digital Signature Standard, in parameter sets ML-DSA-44, ML-DSA-65, and ML-DSA-87 [11]. Private keys occupy 2,560, 4,032, and 4,896 bytes; public keys 1,312, 1,952, and 2,592 bytes; and signatures 2,420, 3,309, and 4,627 bytes [11]. Federal guidance notes that signatures of this magnitude "can strain bandwidth in certain use cases" [3], a constraint that propagates into certificate chains and protocol handshakes rather than remaining a property of the algorithm alone.

FIPS 205 specifies SLH-DSA, the Stateless Hash-Based Digital Signature Standard, across twelve parameter sets built on either SHA2 or SHAKE and offered in small-signature and fast-signing variants [12]. Public keys are 32, 48, or 64 bytes according to security category, while signatures range from 7,856 bytes for SLH-DSA-128s to 49,856 bytes for SLH-DSA-256f [12]. The variant choice is consequential: the fast-signing parameter sets roughly double signature size relative to their small-signature counterparts at equivalent security category [12]. The compensating property is the assumption base. Federal guidance describes SLH-DSA as a "[h]ash-based approach, robust fallback independent of lattice/code assumptions," while conceding "[l]arger signatures (tens of KB), slower signing" [3]. SLH-DSA is therefore best understood not as a competitor to ML-DSA but as insurance against a structural break in lattice cryptography, priced in bandwidth.

Table 2. Key, ciphertext, and signature sizes in bytes for the published module-lattice standards, by parameter set.

Algorithm	Parameter set	Public key	Private key	Ciphertext or signature
ML-KEM	ML-KEM-512	800	1,632	768
ML-KEM	ML-KEM-768	1,184	2,400	1,088
ML-KEM	ML-KEM-1024	1,568	3,168	1,568
ML-DSA	ML-DSA-44	1,312	2,560	2,420
ML-DSA	ML-DSA-65	1,952	4,032	3,309
ML-DSA	ML-DSA-87	2,592	4,896	4,627

Table 2 collects these figures. For ML-KEM, FIPS 203 terms the columns the encapsulation key, decapsulation key, and ciphertext, and ML-KEM additionally produces a 32-byte shared secret key at every parameter set [10]; for ML-DSA they are the public key, private key, and signature [11].

Two further algorithms complete the picture and are routinely misreported. Falcon was selected and "will be published in FIPS 206 (in development)" [1]; **as of August 2026 it is not a published standard, and no initial public draft has been released** — NIST's own project status, given in 2026, records FN-DSA as "[u]nder development" [25]. The algorithm's basis is stated by NIST directly: on publication it "will be dubbed FN-DSA, short for FFT (fast-Fourier transform) over NTRU-Lattice-Based Digital Signature Algorithm" [26], and it follows a hash-and-sign paradigm offering "[s]maller bandwidth and fast verification but more complicated implementation" [25]. HQC was selected for standardization on 11 March 2025 [1] and, **as of August 2026, has no FIPS number and remains at draft stage** — NIST records it as "selected for standardization in 2025 (draft underway)," with "[n]o further KEMs under evaluation" [25]. NIST's stated plan is to issue a draft for public comment and "publish a final version in approximately two years" from March 2025 [4], placing a final HQC standard around 2027. Organisations building migration plans on a set of five available standards are planning against three.

Guidance on correct use has advanced alongside the standards themselves. NIST published SP 800-227, *Recommendations for Key-Encapsulation Mechanisms*, in final form in September 2025, covering the definitions, properties, and applications of KEMs and providing recommendations for implementing and using them securely [13]. The distinction between a specified algorithm and a correctly deployed one is where most implementation risk now resides.

Comparative Trade-offs

Comparing these algorithms to each other, and to what they replace, is less straightforward than tabulating their sizes suggests, and the difficulty is instructive rather than merely technical. FIPS 203 declines the conventional shorthand outright, stating that "security strength is not described by a single number, such as '128 bits of security'" and assigning parameter sets to NIST security categories instead [10]. A comparison table with a bit-strength column would therefore contradict the standard it purports to summarise. The comparison offered here is accordingly built on category alignment, which the standards do define, rather than on equivalent bit strengths, which they decline to assert.

Table 3. Security-category alignment of the standardized and selected post-quantum parameter sets, with the classical key sizes of comparable pre-quantum strength. Units differ by column and are not interchangeable: classical figures are parameter lengths in bits, post-quantum figures are parameter set designations.

NIST security category	Symmetric reference	RSA modulus k (bits)	ECC field f (bits)	ML-KEM	ML-DSA	HQC
— (112-bit strength)	3TDEA	2,048	224–255	—	—	—
1	AES-128	3,072	256–383	ML-KEM-512	— [a]	HQC-1
2	—	—	—	—	ML-DSA-44	—
3	AES-192	7,680	384–511	ML-KEM-768	ML-DSA-65	HQC-3
5	AES-256	15,360	512+	ML-KEM-1024	ML-DSA-87	HQC-5

Table 3 sets out that alignment. Post-quantum category assignments are those claimed by the standards themselves: FIPS 203 states that "ML-KEM-512 is claimed to be in security category 1, ML-KEM-768 is claimed to be in security category 3, and ML-KEM-1024 is claimed to be in security category 5" [10]; FIPS 204 that "ML-DSA-44 is claimed to be in security strength category 2, ML-DSA-65 is claimed to be in category 3, and ML-DSA-87 is claimed to be in category 5" [11]; and the HQC specification assigns HQC-1, HQC-3, and HQC-5 to NIST categories 1, 3, and 5 respectively [27]. Classical values are the comparable strengths tabulated in SP 800-57 Part 1 [28], to which FIPS 203 refers the reader when it states that "[t]he security categories 1-5 are defined in SP 800-57, Part 1" [10].

One qualification governs how the rows should be read. The correspondence between a NIST security category and a particular symmetric key length is established in the evaluation criteria of the standardization process rather than in the published standards themselves, and is reproduced here as the conventional reading rather than quoted from a cited document. What the cited sources establish directly is narrower: that the categories are "defined based on the concrete cost of quantum key search against AES" [9], and that the classical key sizes in each row are of comparable pre-quantum strength to the symmetric reference [28]. The 112-bit row has no post-quantum counterpart, and is included because RSA-2048 and 224-to-255-bit elliptic curves remain widely deployed: the standardized portfolio begins at a strength above much of what it is replacing.

[a] ML-DSA has no parameter set claiming category 1, but ML-DSA-44 falls to that category under the conditions described below.

Two features of this table carry more weight than the alignment itself. The first is that ML-DSA offers no category 1 parameter set, and the gap is not an oversight — the signature standard's smallest parameter set sits a category above the smallest KEM parameter set. The second is that the alignment is asymmetric in what it measures. The classical column reports strengths estimated against classical adversaries, and SP 800-57 attaches an explicit warning to its own figures: "[t]he security-strength estimates will be significantly affected

when quantum computing becomes a practical consideration" [28]. The post-quantum categories, by contrast, are calibrated against quantum adversaries from the outset, being "defined based on the concrete cost of quantum key search against AES" [9]. The two columns are therefore commensurable only in the pre-quantum world that the migration exists to leave. Reading across a row gives the classical key size a post-quantum parameter set replaces; it does not give two measurements of the same quantity.

A further subtlety resists tabulation altogether. Security category is not purely a property of a parameter set but partly of its deployment. FIPS 204 requires that the random bit generator used in ML-DSA key generation "shall have a security strength of at least 192 bits for ML-DSA-65 and 256 bits for ML-DSA-87," while for ML-DSA-44 it "should have a security strength of at least 192 bits and shall have a security strength of at least 128 bits" — and specifies that if a generator with "at least 128 bits of security but less than 192 bits of security is used, then the claimed security strength of ML-DSA-44 is reduced from category 2 to category 1" [11]. An implementer can therefore move a parameter set between categories through a choice made elsewhere in the system. This is a small instance of the paper's larger argument: the security an organisation obtains is a property of its deployment, not of the algorithm named in its procurement documents.

Against this common scale, the portfolio's trade-offs become legible as deliberate design choices rather than a menu. In key establishment, HQC's cost relative to ML-KEM is the price of algorithmic diversification, and it is substantial: at category 1, HQC's encapsulation key is 2,241 bytes against ML-KEM-512's 800, and its ciphertext 4,433 bytes against 768 [27], [10] — roughly a threefold and a near-sixfold increase for equivalent claimed security. Both produce an identical 32-byte shared secret [27], [10], so the entire difference is transmission cost. That figure quantifies what a code-based fallback costs, and it explains why HQC is positioned as insurance against a lattice break rather than as a general-purpose default.

The signature standards trade along a different axis. ML-DSA and SLH-DSA occupy the same categories but rest on different foundations, and the cost of the more conservative assumption is bandwidth: ML-DSA signatures run from 2,420 to 4,627 bytes [11], while SLH-DSA signatures range from 7,856 to 49,856 bytes [12] — an order of magnitude at the upper end. Federal guidance frames the exchange plainly, describing SLH-DSA as a "[h]ash-based approach, robust fallback independent of lattice/code assumptions" while conceding "[l]arger signatures (tens of KB), slower signing" [3]. FN-DSA, when published, is intended to occupy the remaining corner of this space, offering "[s]maller bandwidth and fast verification but more complicated implementation" [25] — and the implementation complexity is not incidental, since it is among the reasons the standard remains in development two years after the others were published.

The practical reading is that no single algorithm dominates. ML-KEM and ML-DSA are the general-purpose defaults on size and speed; SLH-DSA buys assumption diversity with bandwidth; HQC buys mathematical diversity with transmission cost; FN-DSA will buy compactness with implementation difficulty. A portfolio assembled on these terms is only useful to an organisation able to move between its members, which returns the argument to cryptographic agility.

The Algorithms That Were Not Standardized

The candidates eliminated during the process carry more argumentative weight than those selected, because they establish that the evaluation was adversarial and that mathematical confidence is revisable.

Rainbow, one of three finalist signature schemes, fell to classical cryptanalysis. Beullens demonstrated that "given a Rainbow public key for the SL 1 parameters of the second-round submission, our attack returns the corresponding secret key after on average 53 hours (one weekend) of computation time on a standard laptop" [14]. SIKE fell more decisively still: Castryck and Decru's Magma implementation "breaks SIKEp434,

which aims at security level 1, in about ten minutes on a single core" [15]. NIST's own account records that "published cryptanalytic results demonstrated that SIKE was insecure," prompting its removal, and that SIKE's submitters themselves acknowledged its insecurity [4]. Neither break required a quantum computer.

Two remaining candidates were set aside without being broken, and conflating those outcomes with the breaks above misrepresents the record. BIKE was not selected because NIST "does not consider the DFR analysis for BIKE to be as mature as that for HQC," where DFR denotes the decryption failure rate governing IND-CCA2 security [4]. HQC prevailed on the strength of that analysis: "Given the critical need for strong IND-CCA2 security in a general-purpose KEM, HQC was selected for standardization" [4]. Classic McEliece was set aside for reasons that were largely institutional rather than cryptographic. NIST records that "the interest expressed in Classic McEliece was limited, and having more standards to implement adds complexity to protocols and PQC migration," and further that the scheme was "currently under consideration for standardization by the International Organization for Standardization (ISO)," where concurrent standardization "risks the creation of incompatible standards" [4].

Hybrid Constructions

Hybrid deployment combines a classical algorithm with a post-quantum one so that, as federal guidance puts it, "[c]ompromising the security of the operation requires an attacker to break both the classical and the PQC schemes if properly implemented" [3]. In key exchange, the parties generate one traditional and one post-quantum key pair, and the resulting shared secrets are "cryptographically combined (e.g., through the use of a key derivation function taking both as inputs)" [3]. At the protocol level within TLS 1.3, the client signals support by sending key shares "for both a traditional group (e.g., x25519) and a PQC KEM (e.g., FIPS 203 ML-KEM-768)" in the key_share extension, with both secrets fed into the key derivation function [3]. This construction is why TLS 1.3 is treated as a migration prerequisite: agencies must support it "not later than January 2, 2030" [3].

The case for hybrids is not unanimous. The same federal guidance that documents the mechanism cautions that hybrid architecture "introduces its own risks and complexities," that agencies "should perform a thorough evaluation of its tradeoffs," and that it is "an intricate and resource-intensive stopgap" [3]. The eliminations described in the preceding subsection supply the counterargument: two candidates were broken classically during evaluation, so a construction that survives the failure of either component has demonstrable value. The durable conclusion is the one federal guidance itself reaches — cryptographic agility, "an architectural principle that enables an organization to switch its cryptographic algorithms with minimal disruption," which "requires more than simply avoiding hardcoded algorithm names" [3]. Whether a given system deploys a hybrid today matters less than whether it can change algorithms when the next result lands.

8. Migration Mandates and Their Reach

The Coordination Architecture

The United States has assigned the post-quantum transition a governance structure rather than a sponsor. Executive Order 14412 directs that the Director of the Office of Management and Budget and the National Cyber Director, consulting the Assistant to the President for National Security Affairs and the Administrator of the Office of Electronic Government, "lead the strategic coordination and oversight of the national PQC migration policy and strategy," ensuring "its alignment with broader cybersecurity goals" [2]. Technical authority is separated from that oversight function: the Secretary of Commerce, through the Director of NIST

and in consultation with the Director of the National Security Agency and the Director of CISA, is to supply agencies "on an ongoing basis with comprehensive technical guidance on PQC implementation, including best practices in implementation and risk management strategies" [2].

The division is worth noting because it is the structure enterprises are being asked to mirror. Policy ownership, technical standards, and operational assistance sit with different bodies, and the implementing memorandum makes the same separation internally by pairing OMB and the Office of the National Cyber Director for plan review with NIST and CISA for technical support [3]. The memorandum also fixes an important limit on the whole architecture: it "does not apply to national security systems" [3], which are handled instead through an annual report from the Director of the NSA to the President on migration status for agencies operating such systems [2].

Obligations on Federal Agencies

Agency obligations begin with accountability and inventory rather than deployment. Within thirty days each agency head must identify a PQC migration lead [2], and within ninety days OMB must issue guidance requiring agencies to review their inventory of high value assets and high impact systems, excluding national security systems, and to submit a plan for the transition [2]. The definitional scope is narrower than "all federal systems": a high impact system is one in which at least one security objective carries a FIPS 199 potential impact value of "high," while a high value asset is one so designated under OMB Memorandum M-19-03 or its successor [2].

Against that scope the order sets two dates — post-quantum key establishment for all such systems by 31 December 2030, and post-quantum digital signatures by 31 December 2031 [2]. Implementing guidance requires each agency to submit a PQC Migration Plan to OMB and the Office of the National Cyber Director within 120 days [3]. Government also commits to demonstrating the work on itself: NIST must initiate a migration pilot on a subset of its own systems within 180 days, "to be completed no later than December 31, 2027" [2].

Reach Beyond Government: Procurement and Validation

The mandate's significance for private enterprise lies less in what it requires of agencies than in what it propagates through contracting. The Federal Acquisition Regulatory Council is directed, within 180 days, to publish a proposed rule amending the Federal Acquisition Regulation "to require covered contractors to comply by December 31, 2030, with NIST's FIPS, including all applicable FIPS incorporating PQC compliant algorithms" [2]. A second proposed rule, due within 270 days, would extend contractor vulnerability disclosure requirements to "reports of cryptographic vulnerabilities, including testing for lack of encryption and the use of non-FIPS approved algorithms" [2]. Firms that sell to the federal government therefore inherit a 2030 cryptographic deadline through procurement regardless of their own risk assessment.

Supply-side capacity is treated as a constraint on that timetable rather than assumed. Within 180 days NIST must, "to the extent appropriate and consistent with applicable law, revise the processes used by the Cryptographic Module Validation Program to accelerate validations of cryptographic modules" [2]. The inclusion is telling: validated modules, not published algorithms, are what procurement rules actually require, and the queue for validation is a plausible bottleneck between a standard existing and a product being purchasable. Alongside this, the order directs OMB, the Secretary of War, the NASA Administrator, and the Administrator of General Services to identify cost-saving opportunities including "shared procurement of PQC tools, joint training programs, and centralized technical support" [2].

Cloud and the Shared Responsibility Model

Where systems are operated by third parties, obligation must be apportioned before it can be discharged. Agencies are directed to engage their FedRAMP-authorised cloud service providers "to delineate PQC migration responsibilities within the shared responsibility model," while CISA and the Department of War, coordinating with GSA, "will lead PQC migration efforts for FedRAMP-authorized cloud service providers as well as software-as-a-service (SaaS), platform-as-a-service (PaaS), and infrastructure-as-a-service (IaaS) solutions that are used at more than one agency" [3]. Centralising migration for multi-tenant services is an efficiency measure, but it also concentrates schedule risk: a delayed provider delays every tenant. Identity infrastructure receives parallel treatment, with GSA required to establish an inter-agency working group on modernising Federal Identity, Credential, and Access Management to support PQC within sixty days of the memorandum [3].

Critical Infrastructure and Exposure by Data Lifetime

The order extends past federal systems to infrastructure the government does not operate. Agencies serving as Sector Risk Management Agencies under National Security Memorandum 22 are directed to work with CISA "to assist critical infrastructure owners and operators in developing their PQC migration plans" [2], making sector regulators the transmission mechanism into private operators. The effort is also pursued internationally: the Secretary of State, with NIST, DHS, the National Cyber Director, the Secretary of War, and the Director of National Intelligence, is to "identify and engage foreign governments and industry groups in key countries to encourage their transition to PQC algorithms standardized by NIST" [2].

For enterprises outside these channels, the most transferable element of federal policy is its prioritisation criterion. Agencies must rank systems by risk, prioritising high impact systems, high value assets, and any system holding "highly-sensitive data" or judged "likely to be particularly vulnerable to CRQC-based attacks," with particular attention to systems that "[c]ontain data expected to remain mission-sensitive in 2030" [3]. This orders migration by the confidentiality lifetime of the data rather than by the age, cost, or visibility of the system — which is precisely the logic that makes the harvest-now-decrypt-later exposure described in Section 6 actionable. Sectors holding records with multi-decade sensitivity inherit the earliest deadlines under this test, irrespective of whether they are federally regulated.

International Mandates and Points of Divergence

The commitment to engage other governments reflects a landscape in which the United States is one participant rather than the author. Several jurisdictions have published national migration guidance, and reading them together yields two findings that no single roadmap discloses. Each instrument below was retrieved from its own issuing authority; where a national roadmap summarises another country's timeline, that summary was used to locate the primary source rather than cited in its place.

Table 4. National post-quantum migration guidance, by issuing authority. Instruments differ in legal character and the milestones are not equivalent obligations.

Jurisdiction	Instrument	Character	Principal milestones
United States	Exec. Order 14412; OMB M-26-15 [2], [3]	Binding on federal agencies; NSS excluded	Key establishment 2030; digital signatures 2031; full migration 2035
United Kingdom	NCSC, *Timelines for Migration to PQC*, Mar. 2025 [19]	Guidance addressed to all organisations	Discovery 2028; highest-priority activities 2031; complete 2035

Canada	CCCS, ITSM.40.001, Jun. 2025 [20]	Roadmap for federal departments	Departmental plans Apr. 2026; high-priority systems 2031; remaining systems 2035
India	DST Task Force report, Feb. 2026 [17]	National roadmap under the National Quantum Mission	Critical infrastructure: foundations 2027, high-priority migration 2028, full adoption 2029; other enterprises 2028 / 2030 / 2033
European Union	Commission Recommendation (EU) 2024/1101 [21]	Non-binding recommendation	Coordinated Implementation Roadmap within two years of April 2024

The first finding is a convergence on 2035. The United States, the United Kingdom, and Canada independently terminate their schedules in that year, and the American federal timetable and the British guidance also agree on 2031 for high-priority work despite differing in scope and legal force. Convergence of this kind is useful to enterprises operating across jurisdictions, because it means a single migration programme can satisfy several regimes; and it is a reasonable default for organisations in jurisdictions that have published nothing.

India is the outlier, and its divergence runs in the direction of urgency rather than delay. Its Task Force sets accelerated timelines under which critical infrastructure sectors — "defence, power, and telecom" — pursue "Foundations by 2027, High-Priority Migration by 2028, and Full PQC Adoption by 2029," with other enterprises following "the broader timelines of 2028, 2030, and 2033" [17]. Full adoption for critical infrastructure is thus targeted six years ahead of the 2035 horizon adopted elsewhere. The report grounds this schedule in an assessment that the threat window is compressing faster than the consensus allows; that assessment rests substantially on executive commentary rather than technical analysis, and is reported here as the roadmap's stated rationale rather than adopted as a finding of this review. The schedule itself, however, is a published national commitment regardless of the reasoning offered for it, and it establishes that the 2035 convergence is a policy choice rather than an engineering necessity.

The second finding is that these instruments are not comparable obligations, and tabulating their dates without saying so would misrepresent all of them. Executive Order 14412 and OMB M-26-15 bind United States federal agencies and carry procurement consequences. The Canadian roadmap directs federal departments and requires annual progress reporting. British guidance addresses "all organisations" without binding any of them. Commission Recommendation (EU) 2024/1101 is expressly non-binding, and its two-year deadline attaches to producing a roadmap rather than to completing a migration — an instrument one step further removed from operational obligation than the others. A reader comparing "2035" across three columns is comparing a legal duty, an administrative expectation, and a professional recommendation.

Two further characteristics of the Indian programme distinguish it from the others and are worth recording, since it is the most recently published of the roadmaps examined. First, it treats cryptographic agility as a procurement condition rather than an architectural aspiration, directing that common requirements across government tenders "ensure crypto-agile and PQC-compliant assets, along with compulsory Bill of Materials" [17]. This converts into a contractual term what other roadmaps recommend as good practice, and pairs it with a national testing and certification programme whose accredited laboratories are to be operational "by December 2026" [17]. Second, it makes technological sovereignty an explicit objective, mandating "preferential consideration of indigenously developed quantum-safe products and solutions in both public and private organisations, subject to technical suitability and interoperability" [17]. No other roadmap

examined here couples migration to domestic supply in this way, and the tension between sovereign supply and the interoperability that cross-border cryptography requires is unresolved in the document itself.

Sector regulators are beginning to act beneath these national frameworks, though more slowly. The Reserve Bank of India constituted an expert committee in May 2026 on a "Quantum Secure and Adaptive Financial Ecosystem," charged with evaluating "the financial sector's cryptographic inventory through a Cryptography Bill of Materials (CBOM)," conducting "a cross-country analysis and assess[ing] the adequacy of existing regulatory frameworks," and recommending "a roadmap and framework to quantum-secure the Indian financial system," reporting within six months of its first meeting [18]. As of this review the committee has been constituted and its report has not been published; no post-quantum migration requirement has yet been placed on Indian banks. The distinction between a regulator studying the problem and a regulator mandating a response is precisely the one this paper has insisted on for standards, and it applies with equal force to policy.

Coverage here is limited to jurisdictions whose instruments were retrievable in primary form. Australian guidance is reported to set the most aggressive deadline of any national programme — ceasing use of traditional asymmetric cryptography by the end of 2030 — but the primary text could not be retrieved, and the claim is therefore excluded from Table 4 rather than carried on secondary authority. Programmes in Singapore, South Korea, and the United Arab Emirates were identified but not obtained.

9. Deployment Experience: Measured Evidence from Production Systems

Policy establishes obligation; it does not establish feasibility. The evidence that post-quantum cryptography can be operated at scale comes instead from systems already running it, and that evidence has a particular character worth stating before presenting it. Announcements of support are abundant and establish little. What follows is restricted to deployments whose operators have published either a peer-reviewed protocol analysis or measured operational data, and in the latter case the measurements — not the surrounding commentary — are what is cited.

Protocol-Level Deployment and Its Verification

The most thoroughly documented post-quantum deployment is Signal's PQXDH key agreement protocol, which is unusual in having been subjected to formal verification published at a peer-reviewed venue and co-authored by the operator. PQXDH extends the earlier X3DH protocol with a post-quantum key encapsulation mechanism, and its stated purpose is precisely the threat described in the harvest-now-decrypt-later discussion: the protocol "seeks to protect the confidentiality of messages against harvest-now-decrypt-later attacks" [22]. The deployment predates the standard it anticipates — the KEM was instantiated using Kyber, described at the time of analysis as "currently undergoing standardization as ML-KEM" [22] — which is itself instructive. The first large-scale post-quantum messaging deployment shipped against a draft rather than a published FIPS, on the reasoning that the harvest window was open before the standard was final.

The verification result matters more than the deployment announcement. Analysing PQXDH with two formal tools, the authors report that their "analysis identifies several flaws and potential vulnerabilities in the PQXDH specification, although these vulnerabilities are not exploitable in the Signal application, thanks to specific implementation choices" [22]. Two conclusions follow that generalise well beyond Signal. First, the gap between a specification and a sound deployment is where the residual risk in this migration now sits — the same conclusion the key-establishment guidance reaches from the standards side. Second, proving the deployed configuration secure required a property the standards did not supply: the analysis "highlighted the

need for an additional binding property of the KEM, which we formally define and prove for Kyber" [22]. An organisation adopting a standardized primitive inherits the primitive's guarantees, not the guarantees its protocol actually needs.

Measured Adoption, and Where It Stops

Adoption data at internet scale is available from a small number of large infrastructure operators. Cloudflare, which publishes measurements of its own traffic, reported in October 2025 that "over half of human-initiated traffic with Cloudflare is protected against harvest-now/decrypt-later with post-quantum encryption" [23]. Among the top hundred thousand domains measured a month earlier, "39% support PQ now, up from 28% only six months earlier" [23].

The more consequential figure is the one that shows where migration halts. Against more than half of client-facing traffic protected, the share of *origin servers* supporting the same hybrid key agreement was **3.7%** [23]. The asymmetry locates the real state of the transition: post-quantum key agreement has been adopted rapidly where a small number of parties control both endpoints — browsers and content delivery networks, shipping in coordinated releases — and has barely begun in the interior, where migration requires touching heterogeneous, individually administered systems that no single operator can upgrade on a release schedule. This is the same conclusion that the discussion of discovery and legacy estates reaches from policy documents, arrived at independently from measurement, and it suggests that headline adoption percentages describe the tractable half of the problem.

What Deployment Broke

Operational disclosure is more valuable than adoption counting because it records constraints that analysis does not predict. Three are documented.

Protocol ossification produced outright failure. Deploying an early post-quantum key exchange, "a small but significant fraction of clients experienced broken connections with NTRU-HRSS," the cause being "the size of the NTRU-HRSS keyshares" [23]. Larger cryptographic objects broke intermediaries that had assumed the dimensions of the objects they previously carried — a failure mode of the network path rather than of the algorithm.

Size constraints also bound the signature migration that has not yet happened. Certificate chains larger than ten kilobytes are reported to be rejected by "some clients or middleboxes," rendering that approach "a non-starter" [23], and adopting ML-DSA-44 for the web public key infrastructure would add "15kB of data that needs to be transmitted from the server to the client" [23]. Read against the signature sizes in Table 2, this converts an abstract parameter into a deployment constraint, and it explains a pattern visible across every national roadmap examined in the preceding section: key establishment is scheduled first and signatures later, not because signatures matter less, but because the certificate infrastructure cannot yet carry them.

Performance cost, by contrast, proved modest, though the two measurement sources available do not agree on its magnitude and the disagreement is worth stating rather than resolving by selection. The operator disclosure reports that post-quantum key agreement "has already incurred a 4% slowdown in TLS handshake time" [23]. The larger multi-vantage study reaches a stronger conclusion: "[c]ontrary to early experimental studies suggesting measurable overhead, we find that PQ-TLS introduces no meaningful latency increase in Internet settings" [29]. The findings are reconcilable rather than contradictory. They measure different quantities against different baselines — a percentage change in the handshake phase alone, observed by a single operator against its own prior configuration, versus end-to-end connection latency observed across many networks — and a four percent regression confined to the handshake is consistent with no material

effect on what a user experiences. The two agree on what matters for migration planning: computation is not the binding constraint. Size is.

The adoption and failure measurements above come from one large content delivery network and characterise the traffic it observes; the multi-vantage study cited alongside them draws on eleven independent observation points and is used here to check the operator's figures rather than to extend them. Both nevertheless observe web traffic, and both therefore under-represent enterprise interior networks, operational technology, and non-web protocols — the very environments the 3.7% origin figure suggests are furthest behind.

Reported Migration Challenges

Where operator telemetry ends, government assessment provides a complementary account. India's national Task Force, reporting in February 2026, enumerates eight challenges to migration drawn from its consultation across academia, government laboratories, and industry: legacy system complexity, where "diverse and inflexible legacy infrastructures, often lacking crypto-agility, will require redesign or replacement"; interoperability during transition, where the coexistence of classical and quantum-safe cryptography "introduces risks of downgrade or insecure fallback"; vendor readiness gaps; performance and operational impact; a skills shortage reflecting "limited availability of PQC-skilled professionals"; governance and investment continuity; assurance and validation gaps, since "independent validation is critical to ensure correct implementation and prevent reversion to vulnerable cryptography"; and cross-sector coordination risks [17]. The list is notable for being an assessment of a national estate rather than a vendor's characterisation of a market, and for the fact that only two of its eight items are technical properties of the algorithms. The remainder concern organisation, supply chain, assurance, and skills — which is where the following discussion turns.

10. Discussion: Transition Challenges and Organizational Readiness

Technical Integration

The binding technical constraint is rarely the algorithm and frequently the estate it must enter. Federal guidance is candid that inventories have "already identified legacy systems for which migration would be too difficult or costly," and directs that "[s]ystems incapable of supporting PQC or hybrid cryptography must be identified and given priority for replacement or decommissioning" [3]. Migration is thus partly a decommissioning programme, and the guidance's remedy is to absorb the work into existing capital cycles by incorporating "PQC upgrades into planned cloud migrations, software development lifecycles, and hardware-refresh schedules" [3] rather than funding it as a standalone initiative.

Size and performance overheads propagate beyond the cryptographic layer. ML-KEM carries "[l]arger overhead than classical Elliptical Curve Technology (ECC)" [3], with encapsulation keys of 800 to 1,568 bytes and ciphertexts of 768 to 1,568 bytes across its parameter sets [10]. ML-DSA signatures run from 2,420 to 4,627 bytes [11], which federal guidance observes "can strain bandwidth in certain use cases" [3]. SLH-DSA is more demanding still, at 7,856 to 49,856 bytes per signature [12] — characterised as "[l]arger signatures (tens of KB), slower signing" [3]. These figures matter because they land in handshakes and certificate chains, which is also why protocol modernisation is treated as a prerequisite: agencies must support TLS 1.3 or a successor "not later than January 2, 2030" [3].

The durable technical deliverable, however, is not any particular algorithm but the capacity to change it. Cryptographic agility is defined as "an architectural principle that enables an organization to switch its

cryptographic algorithms with minimal disruption," and achieving it "requires more than simply avoiding hardcoded algorithm names" [3], with agencies directed to dedicated NIST guidance on the subject [3]. The concrete measures cited are architectural — modern libraries designed for substitution, such as the OpenSSL 3.x provider model or registration of alternative providers through Java's JCA/JCE, together with configuration-driven selection of algorithms rather than compile-time binding [3].

The case for prioritising agility rests on the standardization record rather than on principle. Two candidates that reached advanced evaluation stages were defeated by classical cryptanalysis: Rainbow's second-round SL 1 parameters yielded a secret key "after on average 53 hours (one weekend) of computation time on a standard laptop" [14], while SIKEp434 fell "in about ten minutes on a single core" [15]. Neither attack required a quantum computer. An organisation that can substitute algorithms cheaply is insured against the next such result; one that cannot is exposed to a class of event that has already occurred twice.

Organizational and Operational Readiness

Where migration obligations are assigned is covered in Section 8; what makes discharging them difficult is an organisational problem before it is a technical one. Implementing guidance is explicit that ownership cannot rest with the security function alone: migration "is *not* only the responsibility of the agency-level Chief Information Officer (CIO) and Chief Information Security Officer (CISO)," but "requires accountability and responsibility for each member of an agency's leadership teams" [3]. Migration touches procurement, application development, and infrastructure refresh simultaneously, so a programme owned solely by security lacks authority over most of its own dependencies.

Discovery is where programmes stall, because an organisation cannot migrate cryptography it cannot locate. Guidance concedes that "manual approaches to discovery and management of cryptography are often insufficient" at scale and directs the use of automation, which "is critical for inventory management, policy enforcement, and compliance reporting" [3]. Third-party components compound the difficulty, since the cryptography inside purchased software is not visible to the purchaser by default. To make it legible, the order requires public guidance on "the minimum elements for a cryptographic bill of materials," enabling "automated assessment of the cryptographic assets utilized by a hardware or software element," within 270 days [2].

Timeline and Roadmap

Federal implementation guidance supplies a five-phase schedule, set out in Table 5, that private-sector roadmaps can be benchmarked against. Phase 1, *Strategy, Planning, and Discovery* (2026–2027), covers "inventory (including HVAs and high impact systems), assessment, strategy definition, awareness and training." Phase 2, *Pilots and Early Migration* (2027–2028), addresses "pilots, executing early migrations of prioritized systems, and refining the migration plan." Phase 3, *Prioritized Migration* (2028–2030), migrates key establishment and requires agencies to "[e]nsure all systems are cryptographically agile." Phase 4, *Signature Migration* (2031), does the same for digital signatures. Phase 5, *Full Migration* (2035), completes "the migration of remaining systems with consideration based on risk assessment and the availability of commercial offerings by 2035" [3]. Two features deserve note: agility appears as a mandated deliverable in two consecutive phases rather than as advice, and the terminal horizon is 2035 — later than the headline 2030 and 2031 obligations imply. Agency plans must align with the NIST transition guidance the memorandum designates as the reference document for migration planning [3].

Table 5. Five-phase migration schedule for United States federal agencies, excluding national security systems [3].

Phase	Name	Years	Required focus

1	Strategy, Planning, and Discovery	2026–2027	Inventory of high value assets and high impact systems; assessment; strategy definition; awareness and training
2	Pilots and Early Migration	2027–2028	Pilots; early migration of prioritised systems; refinement of the migration plan
3	Prioritized Migration	2028–2030	Migrate key establishment for all high value assets and high impact systems; ensure all systems are cryptographically agile
4	Signature Migration	2031	Migrate digital signatures for the same systems; ensure all systems are cryptographically agile
5	Full Migration	2035	Complete migration of remaining systems, based on risk assessment and availability of commercial offerings

Cost

Cost evidence for post-quantum migration is asymmetric, and the distinction governs what can responsibly be claimed. One jurisdiction has published a governmental estimate covering its own civilian systems. No comparable cross-industry or private-sector figure exists. Conflating the two overstates the state of knowledge, because the available number is considerably narrower in scope than its circulation suggests.

The United States estimate appears in the report to Congress required by the Quantum Computing Cybersecurity Preparedness Act. Drawing on cost estimates that agencies submit annually alongside their cryptographic inventories, the Office of the National Cyber Director projects that "the total government-wide cost required to perform a migration of prioritized information systems to PQC between 2025 and 2035 will be approximately $7.1 billion in 2024 dollars" [16]. Three qualifications travel with that figure and should not be detached from it. Its scope is prioritized federal *civilian* systems: national security systems are excluded, with the Department of Defense, the Office of the Director of National Intelligence, and the National Manager for such systems "developing separate funding estimates" [16]. Its window is bounded to 2025–2035. And it is explicitly provisional — the report describes "a high, but expected, level of uncertainty," requires agencies to revise their estimates annually, and states that "[i]nitial cost estimates represent a rough order of magnitude rather than precise calculations" [16]. Quoted without these conditions, as it frequently is, the figure asserts a precision that was never claimed for it.

No equivalent estimate exists for the private sector or across industries, and this review does not supply one. What can be stated is where cost concentrates, and the federal report locates a substantial share of it in equipment that cannot be upgraded at all. Agencies identified systems "whose cryptographic algorithms were hardwired into the hardware or firmware, or those that lack the capacity to accept replacement cryptographic algorithms," and the report records that "[t]he cost to replace those systems constitutes a significant portion of the overall estimate" [16]. This corroborates from the budgetary side what the technical integration discussion argues from the engineering side: migration is in substantial part a replacement programme rather than a reconfiguration, and its cost therefore follows the economics of capital replacement rather than those of software maintenance.

Policy proposes two controls. Executive Order 14412 directs agencies to identify "cost-saving opportunities," naming "migration of cloud-based technologies, shared procurement of PQC tools, joint training programs, and centralized technical support" [2] — an enumeration that identifies procurement, tooling, training, and support as the principal drivers. The complementary control is sequencing rather than spending, absorbing upgrades into planned refresh cycles to "maximize efficiency and minimize costs" [3].

Regulatory and Compliance

The compliance obligations themselves — agency deadlines, contractor requirements, and their statutory basis — are set out in Section 8 and are not restated here. What they impose operationally is a documentation burden distinct from the engineering one: agencies must submit migration plans within 120 days of the implementing memorandum [3], and those plans are expected to be revised rather than filed, being "treated as a dynamic document that will mature over time" [3]. Demonstrating compliance therefore requires sustained evidence of cryptographic inventory and progress, which is why the discovery and automation capabilities described in Section 10.2 are prerequisites for reporting as much as for migration.

11. Specialized Applications: Constrained Devices and Long-Lived Data

Constrained Devices

Post-quantum parameters interact badly with devices whose constraints are fixed at manufacture. The published standards impose costs that are modest on a server and prohibitive on a microcontroller: ML-KEM encapsulation keys of 800 to 1,568 bytes with ciphertexts of 768 to 1,568 bytes [10], ML-DSA signatures of 2,420 to 4,627 bytes [11], and SLH-DSA signatures ranging from 7,856 bytes to 49,856 bytes depending on parameter set and variant [12]. Federal guidance registers the consequence plainly, noting that ML-DSA signatures "can strain bandwidth in certain use cases" and characterising SLH-DSA as carrying "[l]arger signatures (tens of KB), slower signing" [3].

The difficulty is that the algorithm best suited to long-lived embedded roles is also the most expensive. SLH-DSA's appeal for firmware signing and roots of trust is precisely its independence from lattice assumptions — described in federal guidance as a "robust fallback independent of lattice/code assumptions" [3] — yet a signature of tens of kilobytes is a substantial imposition on devices with constrained storage, bandwidth, or power budgets. Where the constraint cannot be met, migration becomes replacement: guidance directs that "[s]ystems incapable of supporting PQC or hybrid cryptography must be identified and given priority for replacement or decommissioning" [3], and that upgrades be absorbed into "hardware-refresh schedules" [3] rather than pursued independently. For device fleets whose service lives are measured in decades, that refresh cycle is the migration schedule.

Long-Lived Data

Data retention lengths convert a future capability into a present exposure. Because adversaries may be "collecting United States information now, and decrypting it later once large-scale quantum computers are operational" [2], the relevant test for archived material is whether its confidentiality requirement outlasts the arrival of a capable machine. Federal policy applies exactly this test operationally, directing agencies to prioritise systems that "[c]ontain data expected to remain mission-sensitive in 2030" [3].

The federal schedule acknowledges that some holdings will not be reached quickly. Its final phase targets 2035 for "the migration of remaining systems with consideration based on risk assessment and the availability of commercial offerings" [3] — nearly a decade after the standards were published. Archives whose sensitivity extends beyond that horizon are, in effect, being protected today by cryptography their custodians already expect to fail.

12. Future Directions

The most immediate open question is the completion of the standard set itself. As of August 2026, Falcon "will be published in FIPS 206 (in development)" and remains unpublished [1], [25], while HQC, selected on 11 March 2025 [1], awaits a draft and a final version that NIST expects "in approximately two years" from that date [4]. Nor is the process closed. NIST has stated that "not all NIST PQC standardization is concluded, as NIST is also currently evaluating additional digital signatures" [4], and the scale of that evaluation is substantial: a call issued in July 2022 drew forty submissions by the June 2023 deadline, of which fourteen advanced to a second round in October 2024 [25]. The selection criterion is diversification rather than performance — NIST is "primarily interested in additional general-purpose signature schemes that are not based on structured lattices," and any lattice candidate "would need to significantly outperform CRYSTALS-Dilithium and FALCON and/or ensure substantial additional security properties" [25]. Key establishment, by contrast, is settled: there is "[n]o on-ramp for KEMs currently planned" [25]. Organisations are therefore migrating onto a signature portfolio that will change during the migration, against a key-establishment portfolio that will not — a condition that argues for treating algorithm substitution as a recurring operation rather than a one-time project, and for expecting the churn on the signature side.

That condition gives cryptographic agility its standing as the durable research and engineering objective. Federal guidance defines it as "an architectural principle that enables an organization to switch its cryptographic algorithms with minimal disruption" and is explicit that achieving it "requires more than simply avoiding hardcoded algorithm names" [3]. What "more" consists of, at the level of protocol design, certificate infrastructure, and hardware roots of trust, remains substantially unresolved and is the natural locus of further work. The standardization record supports the emphasis: two candidates that reached advanced evaluation were defeated classically [14], [15], and a portfolio assembled partly as insurance against such events is only useful to an organisation able to exercise the insurance. Indian national guidance treats the same property as a procurement condition rather than an aspiration, directing that government tenders "ensure crypto-agile and PQC-compliant assets, along with compulsory Bill of Materials" [17] — an approach that converts agility from an architectural preference into a contractual requirement, and whose effectiveness is an open empirical question.

Beyond the completion of the portfolio, four technical problems are open, and they share a common character: the mathematics of the standardized algorithms is settled, while the engineering that carries it is not.

Implementation security is now the active attack surface. With the algorithms fixed, cryptanalysis has moved from the schemes to their realisations, and the current results concern protected implementations rather than naive ones. Recent work analysing "the side-channel vulnerability of masked ML-DSA implementations" demonstrates "how an adversary could launch first, second, and higher-order attacks," contributing a filtering technique "that allows the framework to solve for the secret key from a large number of hints," and instantiates a countermeasure that "mitigates the attacks with the highest noise-tolerance while having very little overhead" [24]. The significant fact is that masking — the standard defence — was the target rather than the answer. This is a materially different posture from the pre-standardization breaks of Rainbow and SIKE: those defeated the mathematics, whereas these defeat particular constructions of a published standard, and the remedy is engineering discipline rather than algorithm replacement. Establishing what constitutes adequate implementation assurance, and how it is certified, is unresolved. National programmes have begun to treat it as infrastructure — India's roadmap commits to operationalising accredited testing laboratories "by December 2026" and to a national testing and certification programme

[17] — but validation regimes for side-channel resistance in post-quantum implementations remain immature.

Protocol-level security properties are not fully inherited from the primitives. Formal verification of a deployed post-quantum key agreement found that proving the deployment secure required "an additional binding property of the KEM," which the authors had to "formally define and prove" themselves [22]. The general problem is that a standardized primitive supplies the guarantees its specification states, which are not necessarily the guarantees a protocol composing it requires. Systematic characterisation of the additional properties that post-quantum primitives must satisfy for common protocol patterns — and their verification for each standardized algorithm — is incomplete, and each gap is discovered at present by analysing deployments one at a time.

The signature migration has no demonstrated path at web scale. Measured deployment shows key establishment succeeding while signature migration remains blocked by size: certificate chains beyond ten kilobytes are already rejected by some intermediaries, and migrating the web public key infrastructure to ML-DSA-44 would add roughly fifteen kilobytes per connection [23]. Every national roadmap examined here schedules signatures after key establishment, which defers rather than solves the problem. Whether the resolution lies in compact signatures from FIPS 206, in restructuring certificate chains, or in changing how trust is conveyed altogether is an open design question with a policy deadline attached to it.

The role of quantum key distribution remains contested and is largely settled in practice. Indian national guidance states the operative position concisely: "PQC remains the most deployable approach, while QKD provides strategic, high-assurance capabilities for specific use-cases" [17]. That formulation — PQC as the general answer, QKD as a niche one — reflects the practical constraint that QKD requires dedicated physical infrastructure. Where such infrastructure is being built for other reasons, hybrid architectures combining both become worth investigating, and India's programme, which pairs an inter-city QKD backbone with PQC at the edges and commits to "establish national testbeds as foundational infrastructure for crypto-agility and hybrid PQC–QKD solutions" [17], is among the few settings in which that combination can be evaluated empirically rather than argued about in principle.

13. Conclusion

The post-quantum transition has passed from anticipation into obligation, and the two developments that moved it are recent and specific. Three Federal Information Processing Standards were published on 13 August 2024 [1], and Executive Order 14412, signed 22 June 2026, attached dates to their adoption: post-quantum key establishment across high value assets and high impact systems by 31 December 2030, digital signatures by 31 December 2031, and contractor compliance with post-quantum FIPS by the earlier of those dates [2].

Precision about what exists matters more in this subject than it usually does, because procurement obligations attach to published standards. Three FIPS are published; Falcon's FIPS 206 remains "in development" [1]; HQC, selected on 11 March 2025 [1], has no FIPS number and no final standard expected before roughly 2027 [4]. The gap between the portfolio as commonly described and the portfolio as currently available is two algorithms wide, and it falls entirely on the signature side.

The standardization record also cautions against confidence in any particular algorithm. Rainbow's second-round SL 1 parameters fell to a laptop in 53 hours [14] and SIKEp434 to a single core in about ten minutes [15], neither attack requiring a quantum computer. The portfolio's diversity is a response to that history, and

guidance accordingly treats cryptographic agility as a mandated deliverable rather than a recommendation — required of United States agencies in two consecutive migration phases [3], and written into Indian procurement as a condition on government tenders [17].

The obligation is also no longer confined to one jurisdiction. National guidance in the United Kingdom [19], Canada [20], and the United States [3] converges on 2035, while India targets full post-quantum adoption for critical infrastructure by 2029 [17] — evidence that the common horizon reflects policy judgement rather than engineering necessity. These instruments are not equivalent: some bind, some advise, and one sets a deadline only for producing a roadmap [21]. Organisations operating across borders should plan against the strictest applicable date rather than the most convenient.

What deployment shows is that obligation and adoption are, so far, loosely coupled. Post-quantum key agreement now protects more than half of client-facing web traffic, yet only a small fraction of origin servers support it [23], [29], and the largest measurement study to date finds "limited correspondence" between national timelines and observed deployment patterns [29]. Migration has advanced fastest where a few parties control both ends of a connection and slowest in the heterogeneous interior that mandates actually address. Deployment has also proved constrained by size rather than computation: key agreement has been absorbed with negligible latency cost, while signature migration remains blocked by certificate sizes that current infrastructure will not carry [23], [29].

Urgency, finally, does not depend on predicting when a capable machine arrives. Because adversaries may collect encrypted material now and decrypt it later [2], exposure is determined by how long data must stay confidential — which is why federal policy prioritises systems holding "data expected to remain mission-sensitive in 2030" [3] rather than forecasting the threat. That test is available to any organisation today, and it is the one this review recommends applying first.

Funding

This research received no external funding.

Conflict of Interest

The authors declare no conflict of interest.

Data Availability Statement

This review analyses no primary datasets. All sources are publicly available documents, cited in full in the references with digital object identifiers or stable uniform resource locators. Standards and policy documents were retrieved from the issuing bodies' public repositories; preprints and conference papers from public archives.

AI Usage Disclosure

The authors used Claude (Anthropic), during preparation of this manuscript for assistance with manuscript organization, structuring and language refinement. All content was reviewed and verified by the authors, who accept full responsibility for it, and no AI tool is listed as an author.

Author Contributions

Conceptualization, Sajeevkrishnan M.A. and Arifa P.; methodology, Arifa P.; investigation and source verification, Arifa P.; validation, Sajeevkrishnan M.A.; writing—original draft, Arifa P.; writing—review and editing, Sajeevkrishnan M.A. and Arifa P.; supervision, Sajeevkrishnan M.A. All authors have read and agreed to the published version of the manuscript.

References

[1] National Institute of Standards and Technology, *Post-Quantum Cryptography Standardization*. Gaithersburg, MD, USA: NIST, 2026. [Online]. Available: https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization

[2] Exec. Order No. 14412, "Securing the Nation Against Advanced Cryptographic Attacks," *Federal Register*, vol. 91, no. 121, pp. 38483–38486, Jun. 25, 2026. [Online]. Available: https://www.federalregister.gov/documents/2026/06/25/2026-12909/securing-the-nation-against-advanced-cryptographic-attacks

[3] Office of Management and Budget, *Execution of the Migration to Post-Quantum Cryptography*, Memorandum M-26-15. Washington, DC, USA: Executive Office of the President, Jun. 24, 2026. [Online]. Available: https://www.whitehouse.gov/wp-content/uploads/2026/06/M-26-15-Execution-of-the-Migration-to-Post-Quantum-Cryptography.pdf

[4] G. Alagic *et al.*, *Status Report on the Fourth Round of the NIST Post-Quantum Cryptography Standardization Process*, NIST IR 8545. Gaithersburg, MD, USA: National Institute of Standards and Technology, Mar. 2025, doi: 10.6028/NIST.IR.8545.

[5] 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, doi: 10.22331/q-2021-04-15-433.

[6] C. Gidney, "How to factor 2048 bit RSA integers with less than a million noisy qubits," arXiv:2505.15917, May 2025, doi: 10.48550/arXiv.2505.15917. (Preprint.)

[7] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM J. Comput.*, vol. 26, no. 5, pp. 1484–1509, Oct. 1997, doi: 10.1137/S0097539795293172.

[8] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proc. 28th Annu. ACM Symp. Theory of Computing (STOC '96)*, 1996, pp. 212–219, doi: 10.1145/237814.237866.

[9] S. Jaques, M. Naehrig, M. Roetteler, and F. Virdia, "Implementing Grover oracles for quantum key search on AES and LowMC," in *Advances in Cryptology – EUROCRYPT 2020*, Lecture Notes in Computer Science. Cham, Switzerland: Springer, 2020, pp. 280–310, doi: 10.1007/978-3-030-45724-2_10.

[10] National Institute of Standards and Technology, *Module-Lattice-Based Key-Encapsulation Mechanism Standard*, FIPS 203. Gaithersburg, MD, USA: NIST, Aug. 2024, doi: 10.6028/NIST.FIPS.203.

[11] National Institute of Standards and Technology, *Module-Lattice-Based Digital Signature Standard*, FIPS 204. Gaithersburg, MD, USA: NIST, Aug. 2024, doi: 10.6028/NIST.FIPS.204.

[12] National Institute of Standards and Technology, *Stateless Hash-Based Digital Signature Standard*, FIPS 205. Gaithersburg, MD, USA: NIST, Aug. 2024, doi: 10.6028/NIST.FIPS.205.

[13] G. Alagic, E. Barker, L. Chen, D. Moody, A. Robinson, H. Silberg, and N. Waller, *Recommendations for Key-Encapsulation Mechanisms*, NIST SP 800-227. Gaithersburg, MD, USA: National Institute of Standards and Technology, Sep. 2025, doi: 10.6028/NIST.SP.800-227.

[14] W. Beullens, "Breaking Rainbow takes a weekend on a laptop," in *Advances in Cryptology – CRYPTO 2022*, Lecture Notes in Computer Science. Cham, Switzerland: Springer Nature, 2022, pp. 464–479, doi: 10.1007/978-3-031-15979-4_16.

[15] W. Castryck and T. Decru, "An efficient key recovery attack on SIDH," in *Advances in Cryptology – EUROCRYPT 2023*, Lecture Notes in Computer Science. Cham, Switzerland: Springer Nature, 2023, pp. 423–447, doi: 10.1007/978-3-031-30589-4_15.

[16] Office of Management and Budget, *Report on Post-Quantum Cryptography*, as required by the Quantum Computing Cybersecurity Preparedness Act, Pub. L. No. 117-260. Washington, DC, USA: Executive Office of the President, Jul. 2024. [Online]. Available: https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/07/REF_PQC-Report_FINAL_Send.pdf

[17] Department of Science and Technology, *Implementation of Quantum Safe Ecosystem in India: Report of the Task Force*. New Delhi, India: Government of India, Feb. 2026. [Online]. Available: https://dst.gov.in/sites/default/files/Report_TaskForce_PQMigration_4Feb26%20(v1).pdf

[18] Reserve Bank of India, *Quantum Secure and Adaptive Financial Ecosystem (Q-SAFE) – Setting Up of an Expert Committee*, Press Release. Mumbai, India: RBI, May 25, 2026. [Online]. Available: https://rbi.org.in/Scripts/BS_PressReleaseDisplay.aspx?prid=62803

[19] National Cyber Security Centre, *Timelines for Migration to Post-Quantum Cryptography*. London, U.K.: NCSC, Mar. 20, 2025. [Online]. Available: https://www.ncsc.gov.uk/guidance/pqc-migration-timelines

[20] Canadian Centre for Cyber Security, *Roadmap for the Migration to Post-Quantum Cryptography for the Government of Canada*, ITSM.40.001. Ottawa, ON, Canada: Communications Security Establishment, Jun. 23, 2025. [Online]. Available: https://www.cyber.gc.ca/en/guidance/roadmap-migration-post-quantum-cryptography-government-canada-itsm40001

[21] Commission Recommendation (EU) 2024/1101 of 11 April 2024 on a Coordinated Implementation Roadmap for the transition to Post-Quantum Cryptography, *Official Journal of the European Union*, OJ L, 2024/1101, Apr. 12, 2024. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202401101

[22] K. Bhargavan, C. Jacomme, F. Kiefer, and R. Schmidt, "Formal verification of the PQXDH post-quantum key agreement protocol for end-to-end secure messaging," in *Proc. 33rd USENIX Security Symp. (USENIX Security '24)*, Philadelphia, PA, USA, 2024.

[23] B. Westerbaan, "State of the post-quantum Internet in 2025," Cloudflare Blog, Oct. 28, 2025. [Online]. Available: https://blog.cloudflare.com/pq-2025/

[24] J. Hermelink, K.-C. Ning, and R. Petri, "Finding and protecting the weakest link," in *Advances in Cryptology – CRYPTO 2025*, Lecture Notes in Computer Science. Cham, Switzerland: Springer Nature Switzerland, 2025, pp. 3–37, doi: 10.1007/978-3-032-01901-1_1.

[25] D. Moody, "The NIST PQC project," presented at the Multi-Party Threshold Cryptography Workshop (MPTS 2026), National Institute of Standards and Technology, Gaithersburg, MD, USA, 2026. [Online]. Available: https://csrc.nist.gov/csrc/media/presentations/2026/mpts2026-3b1/images-media/mpts2026-3b1-slides-nist-pqc-moody.pdf

[26] National Institute of Standards and Technology, *NIST Releases First 3 Finalized Post-Quantum Encryption Standards*. Gaithersburg, MD, USA: NIST, Aug. 13, 2024. [Online]. Available: https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards

[27] P. Gaborit *et al.*, *Hamming Quasi-Cyclic (HQC)*, specification, Aug. 22, 2025. [Online]. Available: https://pqc-hqc.org/doc/hqc_specifications_2025_08_22.pdf

[28] E. Barker, *Recommendation for Key Management: Part 1 – General*, NIST SP 800-57 Part 1 Rev. 5. Gaithersburg, MD, USA: National Institute of Standards and Technology, May 2020, doi: 10.6028/NIST.SP.800-57pt1r5.

[29] N. Wickramasinghe, F. Li, S. Jha, and A. Shaghaghi, "Mind the gap: Policy vs reality in post-quantum TLS deployment," in *Proc. ACM Internet Measurement Conf. (IMC '26)*, 2026. [Online]. Available: https://arxiv.org/abs/2607.29005

[30] S. Riou, J.-Y. Park, L. Anwar, A. Poschmann, and M. Hutter, "Apples, oranges, and signatures: Pitfalls and methodology in ML-DSA benchmarking," Cryptology ePrint Archive, Paper 2026/1333, 2026. [Online]. Available: https://eprint.iacr.org/2026/1333

[31] V. Mavroeidis, K. Vishi, M. D. Zych, and A. Jøsang, "The impact of quantum computing on present cryptography," *Int. J. Adv. Comput. Sci. Appl.*, vol. 9, no. 3, pp. 405–414, Mar. 2018. [Online]. Available: https://arxiv.org/abs/1804.00200

[32] R. Acharya *et al.* (Google Quantum AI), "Quantum error correction below the surface code threshold," *Nature*, vol. 638, pp. 920–926, 2025, doi: 10.1038/s41586-024-08449-y.