ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1288

# Quantum Computing and the Future of Cybersecurity: Assessing the Threat to Classical Encryption Standards

Ramya[1]; Rashmi[2]

[1] Mahatma Gandhi Memorial College Udupi
[2] Mahatma Gandhi Memorial College Udupi

Publication Date: 2026/07/31

**Abstract: The maturation of quantum computing hardware and algorithms has transformed what was once a theoretical concern into an active planning priority for governments, enterprises, and standards bodies. This paper assesses the current and projected threat that quantum computation poses to classical public-key cryptographic standards, namely RSA, elliptic curve cryptography (ECC), and Diffie-Hellman key exchange, all of which derive their security from mathematical problems that Shor's algorithm renders tractable on a sufficiently large fault-tolerant quantum computer. We synthesize recent hardware progress, algorithmic refinements that have substantially reduced estimated qubit requirements, and the finalization of NIST's post-quantum cryptography (PQC) standards (FIPS 203, 204, and 205) to provide an evidence-based assessment of where the field stands as of mid-2026. Particular attention is given to the widening gap between optimistic resource estimates, including a contested March 2026 claim of a thousand-fold reduction in factoring resources, and the practical reality that no existing quantum computer approaches the logical qubit counts required to threaten deployed encryption. We further examine migration challenges, including performance overhead, cryptographic agility, constraints on embedded and IoT systems, and the harvest-now-decrypt-later threat model that makes migration urgency independent of exact timeline predictions. Drawing on NIST transition guidance and recent G7 coordination targets, we argue that organizations should treat cryptographic migration as a present-tense operational requirement rather than a future contingency, regardless of unresolved disagreement among experts about when a cryptographically relevant quantum computer will exist. The paper concludes with practical recommendations for risk-based migration prioritization and identifies open research questions in cryptographic agility and quantum-resistant system design.**

***Keywords:*** *Post-Quantum Cryptography, Quantum Computing, Shor's Algorithm, RSA, Elliptic Curve Cryptography, ML-KEM, ML-DSA, NIST, Cryptographic Migration, Harvest-Now-Decrypt-Later.*

**How to Cite:** Ramya; Rashmi (2026) Quantum Computing and the Future of Cybersecurity: Assessing the Threat to Classical Encryption Standards. *International Journal of Innovative Science and Research Technology*, 11(7), 2206-2215. https://doi.org/10.38124/ijisrt/26jul1288

## I. INTRODUCTION

Modern digital infrastructure rests on a small number of mathematical assumptions. The security of online banking, encrypted messaging, virtual private networks, code-signing systems, and the public key infrastructure underlying the World Wide Web all depend on the computational difficulty of integer factorization and the discrete logarithm problem, including its elliptic curve variant. These assumptions have held since the 1970s not because anyone has proven that no efficient classical algorithm exists for these problems, but because decades of sustained cryptanalytic effort have failed to find one. Quantum computing threatens to invalidate this assumption entirely, not through incremental improvement but through a qualitatively different mode of computation for which an efficient algorithm is already known.

Peter Shor's 1994 discovery that a quantum computer could factor large integers and compute discrete logarithms in polynomial time established, in principle, that a sufficiently large and reliable quantum computer would render RSA, Diffie-Hellman, and elliptic curve cryptography insecure. For three decades this remained a largely theoretical concern, since the physical qubit counts and error correction overhead required to run Shor's algorithm against cryptographically meaningful key sizes were many orders of magnitude beyond anything achievable in the laboratory. That comfortable distance has been narrowing. Between 2025 and early 2026, a sequence of algorithmic refinements reduced estimated physical qubit requirements for breaking RSA-2048 from roughly twenty million to under one million, and a March 2026 Google Quantum AI study suggested that elliptic curve cryptography of the kind used in many financial systems

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1288

could be vulnerable to a machine with fewer than half a million physical qubits, an estimate roughly twenty times smaller than prior projections [1], [2]. Separately, a March 2026 academic preprint exploring neutral-atom architectures proposed designs requiring as few as ten to twenty thousand atomic qubits to implement a working version of Shor's algorithm against certain targets [3].

These developments matter independent of whether a cryptographically relevant quantum computer (CRQC) is five, fifteen, or thirty years away, because of a structural feature of the threat: adversaries can record encrypted traffic today and decrypt it retroactively once sufficient quantum capability exists. This harvest-now-decrypt-later strategy means that any information requiring confidentiality beyond the eventual arrival of a CRQC, including state secrets, long-lived intellectual property, genomic data, and financial records, is already exposed under current encryption practices. The urgency of migration to quantum-resistant cryptography is therefore decoupled from precise timeline forecasting; it depends instead on the shelf life of the data being protected today.

In response, the United States National Institute of Standards and Technology (NIST) finalized the first three post-quantum cryptography standards in August 2024 after an eight-year, multi-round public evaluation process: FIPS 203 (ML-KEM, a key encapsulation mechanism derived from CRYSTALS-Kyber), FIPS 204 (ML-DSA, a digital signature scheme derived from CRYSTALS-Dilithium), and FIPS 205 (SLH-DSA, a hash-based signature scheme derived from SPHINCS+) [4], [5]. These standards represent the most concrete artifact of the field's response to the quantum threat, but their existence does not by itself solve the migration problem. Organizations must still inventory cryptographic dependencies across legacy systems, absorb performance and memory overheads that can be substantial on constrained hardware, redesign protocols for cryptographic agility, and coordinate migration across supply chains that may include components with operational lifespans of decades.

This paper makes four contributions. First, it provides a current, evidence-based account of the quantum threat to classical encryption that explicitly distinguishes confirmed hardware and algorithmic progress from contested or unverified claims, using the disputed March 2026 "JVG algorithm" episode as a case study in how such claims should be evaluated. Second, it offers a structured technical overview of the finalized NIST PQC standards, including their security foundations, parameter sets, and comparative performance characteristics. Third, it synthesizes the practical migration challenges facing organizations, drawing on NIST's own transition guidance and recent benchmarking data on embedded and resource-constrained systems. Fourth, it situates the technical material within the current policy landscape, including the G7 Cyber Expert Group's January 2026 coordination statement, to argue for a risk-prioritized, present-tense approach to migration that does not wait for timeline certainty.

The remainder of this paper is organized as follows. Section II establishes theoretical background on classical public-key cryptography and the quantum algorithms that threaten it. Section III examines the evolving threat landscape, including hardware progress, algorithmic improvements, and the harvest-now-decrypt-later model. Section IV provides a detailed technical treatment of the NIST PQC standards. Section V analyzes migration challenges. Section VI reviews the policy and governance response. Section VII discusses synthesis, recommendations, and limitations, and Section VIII concludes.

### *A. Related Work and Positioning*

A substantial body of prior work has addressed individual components of the quantum threat to cryptography: theoretical treatments of Shor's and Grover's algorithms, NIST's own standardization documentation, and a growing applied literature benchmarking post-quantum algorithm performance on specific hardware platforms. This paper's contribution relative to that literature is primarily synthetic and evaluative rather than the introduction of new cryptanalytic or algorithmic results. Where prior surveys have tended either to focus narrowly on the cryptographic mechanics of the NIST standards or to address migration logistics in isolation from the underlying threat model, this paper integrates three strands that are too often treated separately: the rapidly evolving and frequently contested evidence base on quantum hardware and algorithmic threat capability; the finalized technical specification of the post-quantum standards intended to address that threat; and the organizational, economic, and policy dimensions of migration, which ultimately determine whether the available standards translate into deployed protection. A further point of methodological emphasis, motivated directly by the March 2026 JVG episode discussed in Section III-D, is the explicit treatment of claim provenance and verification status as a first-class consideration in threat assessment, rather than treating all published qubit-requirement estimates as equally reliable regardless of their review pathway.

### *B. Methodology*

This paper adopts a structured narrative literature review methodology appropriate to a fast-moving technical and policy domain in which formal standards documents, peer-reviewed and preprint research, and primary-source policy statements must be synthesized alongside contemporaneous reporting on contested claims. Sources were selected to prioritize, in descending order of evidentiary weight, finalized standards documents and government technical reports; peer-reviewed publications and preprints from established research groups; primary statements from major research organizations; and journalistic or expert-commentary sources used principally to establish chronology, public reception, and expert survey data rather than as the basis for technical claims. Where a claim's verification status was contested, as with the JVG algorithm, this paper reports the claim, the basis for the dispute, and the resulting scientific consensus, rather than presenting the original claim without qualification. This approach is intended to model the kind of source-criticism that practitioners and researchers should apply when monitoring this domain going forward, given the

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1288

demonstrated speed at which both genuine algorithmic progress and unsubstantiated capability claims now propagate through technical and public discourse.

## II. BACKGROUND AND THEORETICAL FOUNDATIONS

### *A. Classical Public-Key Cryptography*

Public-key cryptography, introduced by Diffie and Hellman and subsequently realized in deployable form through the RSA cryptosystem, solves the key distribution problem that plagued symmetric cryptography by allowing two parties to establish secure communication without a previously shared secret. RSA's security rests on the assumed difficulty of factoring the product of two large prime numbers; given a public modulus that is the product of two primes each several hundred digits long, no known classical algorithm can recover the prime factors in time polynomial in the size of the modulus. Elliptic curve cryptography achieves comparable security with substantially smaller key sizes by relying instead on the elliptic curve discrete logarithm problem, the difficulty of determining how many times a known base point on an elliptic curve must be added to itself to reach another known point on the curve. Both RSA and ECC, along with Diffie-Hellman key exchange in its classical and elliptic curve forms, belong to a class of cryptographic problems sometimes described as having "structure": they are built on number-theoretic constructions for which the relevant hard problem fits a pattern, the hidden subgroup problem, that quantum computers are unusually well suited to solve.

This structural property is the crux of the quantum threat. Symmetric cryptography, including AES, does not rely on this kind of algebraic structure and is consequently far less affected by quantum algorithms; the principal quantum threat to symmetric ciphers comes from Grover's algorithm, which provides a quadratic speedup for unstructured search and effectively halves the bit-security of a symmetric key, a manageable problem solved by doubling key lengths. Public-key cryptography faces an entirely different category of threat because Shor's algorithm provides an exponential, not merely quadratic, speedup against the specific mathematical problems underlying RSA, Diffie-Hellman, and ECC.

It is worth dwelling on why this asymmetry exists, since it is frequently misunderstood in non-specialist treatments of the quantum threat. A common misconception holds that quantum computers are simply faster classical computers and will therefore eventually break every cryptographic scheme given enough engineering progress. This is incorrect. Quantum speedups arise only for problems that can be cast in a form quantum algorithms are suited to exploit, principally periodicity-finding and certain search structures, and the vast majority of computational problems, including most NP-hard problems believed to underlie symmetric cryptography and many post-quantum candidate schemes, show no known exponential quantum speedup. RSA, Diffie-Hellman, and ECC are exceptional specifically because the hidden subgroup problem, which the quantum Fourier transform solves efficiently, captures their underlying hard problems exactly. Lattice problems, code-based problems, and hash-function-based constructions, the foundations of the NIST post-quantum standards examined in Section IV, do not currently reduce to the hidden subgroup problem in a way that would let Shor's algorithm or a known variant apply, which is precisely why they were selected as quantum-resistant replacements.

### *B. Quantum Computing Principles Relevant to Cryptanalysis*

A quantum computer manipulates information encoded in quantum bits, or qubits, which unlike classical bits can exist in superpositions of the 0 and 1 states and can be entangled with one another such that the state of one qubit is correlated with the state of another regardless of physical separation. Quantum algorithms exploit superposition and entanglement to perform certain computations that would require exponential time classically in polynomial time, but only for problems with the right kind of mathematical structure; quantum computers do not provide a general speedup for all computational problems, and for many tasks they offer no advantage at all.

Physical qubits are noisy and decohere rapidly, which means that any computation of practical cryptographic significance requires quantum error correction, encoding many physical qubits into a smaller number of more reliable logical qubits. The overhead of error correction is the central reason why the gap between today's quantum hardware, which fields machines with at most a few dozen logical qubits, and the thousands of logical qubits required to run Shor's algorithm against RSA-2048 remains so large [6]. Reported "physical qubit" counts in vendor announcements should not be confused with the logical qubit counts relevant to cryptographic threat assessment, a distinction that is frequently elided in popular reporting on quantum computing milestones.

### *C. Shor's Algorithm and Grover's Algorithm*

Shor's algorithm solves integer factorization and the discrete logarithm problem by reducing each to the problem of finding the period of a function, which a quantum computer can determine efficiently using the quantum Fourier transform. Running Shor's algorithm against a 2048-bit RSA modulus requires several thousand logical qubits and a circuit depth, the number of sequential quantum operations, sufficient to execute modular exponentiation reliably before the quantum state decoheres. Industry and academic estimates of the physical qubit count required to support this logical qubit count, after accounting for error correction overhead, have fallen substantially over the past several years. Early estimates from the mid-2010s placed the figure near twenty million physical qubits for RSA-2048; algorithmic refinements published by Gidney and others in 2025 reduced this to under one million physical qubits under comparable assumptions, primarily by improving the efficiency of the modular arithmetic circuits and error correction codes used in the construction [1], [7].

Grover's algorithm, by contrast, provides a quadratic speedup for searching an unstructured space, which translates into halving the effective key length of symmetric ciphers and hash functions. A 128-bit AES key, for instance, would offer

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1288

only 64 bits of security against a quantum adversary using Grover's algorithm, still computationally infeasible to brute-force, which is why NIST's post-quantum migration guidance treats symmetric cryptography as comparatively low risk and instead recommends straightforward key-length increases, such as moving to AES-256, rather than wholesale algorithm replacement [8].

## III. THE EVOLVING THREAT LANDSCAPE

### *A. Hardware Progress and Qubit Requirement Estimates*

As of mid-2026, the largest publicly disclosed quantum computers field on the order of dozens of logical qubits after error correction, several orders of magnitude short of the thousands of logical qubits required to threaten RSA-2048 or comparable ECC key sizes [9]. This gap is the single most important fact governing near-term risk assessment: no quantum computer in existence or publicly announced is close to cryptographically relevant. However, the trajectory of resource estimates has been moving in a consistent direction, downward, which is the development that has reshaped expert risk assessments over the past two years.

The physical-qubit estimate for breaking RSA-2048 has fallen from approximately twenty million to under one million as a result of algorithmic and error-correction improvements rather than raw hardware scaling [1], [7]. This distinction matters: progress toward the quantum threat comes from two largely independent channels, the brute engineering challenge of building and maintaining more physical qubits, and the comparatively cheaper intellectual work of designing more efficient algorithms and error-correcting codes that reduce the resource cost of a given attack. The latter channel has produced faster gains recently than the former, which is part of why timeline estimates have compressed even though no headline-grabbing hardware breakthrough has occurred.

Table 1 summarizes the progression of resource estimates discussed in this section, distinguishing verified, peer-reviewed or preprint-published results from the disputed claim examined in Section III-D. Reading the table chronologically illustrates the consistent downward trend in estimated attack cost, punctuated by a single outlier claim that did not survive scientific scrutiny.

Table 1. Evolution of Estimated Physical Qubit Requirements for Breaking Public-Key Cryptography

| Estimate source and year | Target | Estimated physical qubits | Verification status |
|---|---|---|---|
| Early estimates, mid-2010s | RSA-2048 | ~20,000,000 | Established baseline |
| Gidney, 2025 | RSA-2048 | <1,000,000 | Peer-reviewed/preprint, accepted |
| Google Quantum AI, Mar. 2026 | ECC (secp256k1-class) | <500,000 (~1,200 logical qubits) | Published study, engaged constructively by field |
| Caltech-Berkeley-Oratomic, Mar. 2026 | ECC (Bitcoin-class) / RSA-2048 | ~26,000 (ECC) / substantially more (RSA) | Theoretical preprint, design proposal |
| AQTI "JVG algorithm," Mar. 2026 | RSA-2048 / ECC | <5,000 | Press release; disputed and not substantiated under scrutiny [13] |

The contrast between the third and fifth rows of Table 1 is instructive: both were announced within the same month, yet one originated through standard preprint and research-community channels and has been treated as a credible, if still theoretical, contribution to the field, while the other originated through a press release and was rapidly and publicly disputed by leading researchers. This contrast underscores the methodological point developed further in Section III-D, that the provenance and review pathway of a capability claim is as important to threat assessment as the claim's headline numbers.

### *B. Algorithmic Improvements Reducing Resource Requirements*

In March 2026, Google's Quantum AI team published a study showing an approximately twenty-fold reduction in the physical qubits required to solve the mathematical problem underlying elliptic curve cryptography, estimating that a quantum computer with fewer than half a million physical qubits and roughly 1,200 logical qubits could threaten curves such as secp256k1, used in Bitcoin, Ethereum, and a range of financial and identity systems [2], [10]. This result is notable not because it suggests an imminent attack, current hardware remains far short of even this reduced threshold, but because it illustrates how quickly resource estimates can shift through algorithmic insight alone, without any corresponding change in physical hardware capability.

A separate March 2026 preprint from a Caltech, Berkeley, and Oratomic collaboration examined neutral-atom quantum computing architectures and proposed designs in which Shor's algorithm could in principle be implemented with as few as ten to twenty thousand atomic qubits for certain targets, with one proposed configuration of approximately 26,000 qubits estimated to threaten Bitcoin-style elliptic curve cryptography within days, while factoring a full 2048-bit RSA modulus was estimated to require substantially greater resources and time [3], [11]. These figures remain theoretical design proposals rather than demonstrated capabilities, and translating an architectural proposal into a working machine involves substantial additional engineering risk that historically has caused real-world timelines to lag theoretical proposals by years.

### *C. The Harvest-Now-Decrypt-Later Threat Model*

The harvest-now-decrypt-later threat model describes an adversary strategy in which encrypted data is intercepted and stored today, with decryption deferred until a

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1288

cryptographically relevant quantum computer becomes available. This strategy is attractive to well-resourced state and criminal actors because it requires no quantum capability at the time of collection, only sufficient storage and the expectation that decryption will eventually become feasible. The threat model has a direct and important consequence for migration planning: the relevant question for any given dataset is not "when will quantum computers be able to break current encryption" but "will this data still need to be confidential after that date." For data with multi-decade confidentiality requirements, such as national security material, genomic and health records, intellectual property tied to long product lifecycles, and infrastructure design documents, the harvest-now-decrypt-later model means that today's encryption choices are already determining tomorrow's exposure, regardless of how the precise CRQC timeline ultimately resolves.

*D. Case Study: The JVG Algorithm Controversy (March 2026)*

In early March 2026, the Advanced Quantum Technologies Institute (AQTI) issued a press release describing a new factoring algorithm, designated JVG after its credited authors, claiming a thousand-fold reduction in the quantum resources required to break RSA encryption and asserting that the algorithm could factor RSA-2048-class moduli using fewer than 5,000 qubits in approximately eleven hours [12], [13]. The claim received rapid uptake in technology and mainstream media, with some coverage suggesting that internet-scale encryption infrastructure was on the verge of imminent failure.

The scientific community's response was swift and largely critical. Prominent quantum computing researchers, including Scott Aaronson, publicly disputed the claim's technical validity, and the announcement was found to have originated from a press release rather than a peer-reviewed publication, with the underlying technical claims not withstanding independent scrutiny [13]. This episode is instructive for two reasons that bear directly on how this paper, and the field generally, should evaluate threat claims. First, it illustrates the asymmetry between the speed at which a dramatic capability claim can propagate through media channels and the slower, more rigorous process of peer review and independent replication that determines whether such claims hold up. Second, it demonstrates that the cryptographic research and security communities now respond quickly and critically to unverified claims, a sign of a maturing field rather than one prone to uncritical alarm. For the purposes of risk assessment, the JVG episode should be weighted as a contested and currently unsubstantiated claim, distinct from the verified algorithmic improvements described in Section III-B, which were published through standard peer-reviewed or preprint channels and have been engaged with constructively by the research community rather than disputed on technical grounds.

*E. Expert Timeline Estimates and Survey Data*

Despite the compression of resource estimates described above, structured expert surveys continue to place the median estimate for a cryptographically relevant quantum computer in the 2030 to 2035 range, with the annual Global Risk Institute survey conducted by cryptographer Michele Mosca consistently identifying this window as carrying the greatest weight of expert probability mass, while acknowledging a meaningful minority view that places the timeline considerably later or, in some cases, expresses skepticism that a CRQC of the required scale will be built within any specified horizon [9], [14]. A separate informal estimate following the March 2026 algorithmic developments suggested a narrower, more aggressive scenario in which quantum computers capable of credibly threatening 1024-bit RSA in research conditions could emerge as early as 2028 to 2030, though this estimate should be read as one practitioner's extrapolation rather than a consensus position [15]. Google's own public statements in 2026 referenced the possibility of certain encrypted systems becoming vulnerable by 2029, a timeline narrower than many prior expert consensus estimates, though this statement concerned a specific resource threshold rather than a confirmed engineering roadmap for reaching it [16].

The appropriate synthesis of this evidence is neither alarmist nor dismissive. No current or publicly announced quantum computer is within several orders of magnitude of cryptographic relevance. At the same time, the consistent downward trajectory of resource estimates over the past two years, achieved primarily through algorithmic rather than hardware progress, means that the lower end of expert timeline distributions has become more plausible than it appeared as recently as 2023 or 2024. This is the central reason that standards bodies and governments have moved from treating post-quantum migration as a longer-term research priority to treating it as an active operational requirement, a shift documented in the policy responses discussed in Section VI.

## IV. POST-QUANTUM CRYPTOGRAPHIC STANDARDS

*A. NIST Standardization Process*

NIST initiated its post-quantum cryptography standardization effort in 2016, soliciting algorithm proposals from the global cryptographic research community and subjecting submissions to multiple rounds of public cryptanalysis over a six-year evaluation period before announcing its initial selections in July 2022 [4]. This open, multi-round process, modeled on the approach NIST used for the Advanced Encryption Standard competition in the late 1990s, is intended to maximize confidence that selected algorithms have withstood sustained adversarial scrutiny from the research community before being deployed at global scale. The finalized standards were published on August 13, 2024, concluding the eight-year effort, with three Federal Information Processing Standards (FIPS) issued simultaneously [5], [17].

*B. FIPS 203 (ML-KEM)*

FIPS 203 specifies the Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM), derived from the CRYSTALS-Kyber submission, and is intended as the primary general-purpose standard for establishing a shared

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1288

secret key between two parties over a public channel [18]. ML-KEM's security rests on the conjectured hardness of the Module Learning With Errors (Module-LWE) problem, a lattice-based problem believed to resist both classical and quantum cryptanalysis. ML-KEM is offered in three parameter sets, ML-KEM-512, ML-KEM-768, and ML-KEM-1024, corresponding to NIST security categories 1, 3, and 5 respectively, with public keys ranging from approximately 800 to 1,568 bytes and ciphertexts in a comparable range [5]. NIST and independent commentators have noted ML-KEM's relatively small key sizes and competitive computational speed as key advantages favoring its selection as the default general-purpose mechanism [18], [19].

*C. FIPS 204 (ML-DSA)*

FIPS 204 specifies the Module-Lattice-Based Digital Signature Algorithm (ML-DSA), derived from CRYSTALS-Dilithium, as the primary standard for post-quantum digital signatures, addressing use cases including code signing, certificate issuance, and authentication protocols [5], [19]. ML-DSA relies on both the Module-LWE and Module Short Integer Solution (Module-SIS) problems and produces signatures ranging from approximately 2,420 to 4,595 bytes across its parameter sets, mapped to security categories analogous to AES-128, AES-192, and AES-256 equivalence once Grover's algorithm's quadratic speedup against symmetric primitives is accounted for [5]. Industry commentary suggests ML-DSA is expected to serve the substantial majority of digital signature use cases in the post-quantum environment, given its favorable balance of signature size and computational performance relative to the available alternatives [19].

*D. FIPS 205 (SLH-DSA)*

FIPS 205 specifies the Stateless Hash-Based Digital Signature Algorithm (SLH-DSA), derived from SPHINCS+, intended as a structurally independent backup to ML-DSA [5], [19]. Where ML-DSA and ML-KEM both rely on lattice-based hardness assumptions, and would both be undermined by any future cryptanalytic breakthrough against lattice problems, SLH-DSA's security rests solely on the collision resistance of cryptographic hash functions, a substantially more conservative and long-studied assumption. This comes at a considerable cost: SLH-DSA signatures range from approximately 7,856 to 49,856 bytes, one to two orders of magnitude larger than ML-DSA signatures, and the scheme is markedly slower in both signing and verification [5]. SLH-DSA's intended role is as an insurance policy rather than a primary mechanism, providing a fallback whose security does not depend on the same mathematical assumptions as the other standardized algorithms, in case future research reveals an unexpected weakness in lattice-based cryptography.

*E. Future Algorithms (FN-DSA, HQC, and the Fourth-Round Process)*

A fourth algorithm, FN-DSA (FFT over NTRU-Lattice-Based Digital Signature Algorithm, derived from FALCON), has been selected for future standardization but had not been published as a finalized FIPS as of the period covered by this study, with NIST having indicated plans to complete this work after the initial three standards [5], [20]. NIST has additionally continued a fourth round of evaluation for alternative key-encapsulation mechanisms, including HQC (Hamming Quasi-Cyclic), which is based on code-based cryptography rather than lattice problems and is intended to provide algorithmic diversity analogous to the relationship between SLH-DSA and ML-DSA, in case future cryptanalysis reveals weaknesses specific to lattice-based assumptions [20], [21]. NIST has explicitly stated its intention to select one or more additional KEM alternatives to ML-KEM in the future, reflecting a deliberate strategy of not placing all post-quantum security on a single mathematical foundation [21].

*F. Comparative Summary*

Table 2 consolidates the technical characteristics of the three finalized standards discussed above, providing a single reference point for the comparative discussion of migration overhead in Section V. The table highlights the central tradeoff facing implementers: ML-KEM and ML-DSA offer compact keys and competitive performance at the cost of relying on a relatively young mathematical assumption, lattice hardness, while SLH-DSA offers a much longer-studied security foundation at substantial cost in signature size and computational speed.

Table 2. Comparative Characteristics of Finalized NIST Post-Quantum Standards

| Standard | Algorithm | Primary use | Hardness assumption | Key/signature size range | Relative performance |
|---|---|---|---|---|---|
| FIPS 203 | ML-KEM | Key encapsulation | Module-LWE (lattice) | ~800-1,568 bytes (keys) | High; competitive with classical ECC |
| FIPS 204 | ML-DSA | Digital signatures | Module-LWE / Module-SIS (lattice) | ~2,420-4,595 bytes (signatures) | High; default choice for most use cases |
| FIPS 205 | SLH-DSA | Digital signatures (backup) | Hash-function collision resistance | ~7,856-49,856 bytes (signatures) | Low; one to two orders of magnitude slower |

As Table 2 indicates, the practical default for most organizations migrating digital signature infrastructure will be ML-DSA, with SLH-DSA reserved for contexts where the structurally independent hardness assumption is judged worth the substantial size and performance penalty, such as long-term archival signatures or root-of-trust certificates whose validity must be assured even against a hypothetical future break of lattice-based cryptography.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1288

## V. MIGRATION CHALLENGES

### *A. Performance and Resource Overheads*

Empirical benchmarking of the finalized PQC standards reveals substantial and unevenly distributed performance costs relative to classical algorithms. On general-purpose server hardware such as the Intel i7-12700K, ML-KEM-768 has been benchmarked at roughly 85,000 operations per second, and on data-center-class GPU hardware such as the NVIDIA H100, throughput reaches into the millions of operations per second, figures that pose little practical obstacle for most server-side deployments [22]. The picture changes considerably on constrained hardware: the same ML-KEM-768 operation has been benchmarked at approximately 150 operations per second on an ARM Cortex-A53 processor, a chip class common in embedded systems, IoT devices, and smart cards, a difference of roughly three orders of magnitude between the most and least capable platforms tested [22]. This range means that algorithm and parameter selection cannot be made uniformly across an organization's infrastructure; the appropriate choice depends heavily on the deployment context, and a single default configuration suitable for cloud infrastructure may be entirely impractical for an embedded sensor network.

Energy consumption presents a related concern for battery-powered and energy-constrained deployments. Reported figures indicate that post-quantum cryptographic operations increase energy consumption by approximately 15 to 50 percent on embedded systems relative to classical equivalents, with the precise figure depending on the choice of algorithm, ML-KEM against SLH-DSA, and the underlying hardware architecture, ARM against x86 [22]. Memory requirements compound this challenge, with post-quantum algorithms requiring two to five times the memory footprint of their classical counterparts, a constraint that is often more binding than raw computational throughput for devices with long operational lifecycles and limited opportunity for hardware replacement, including industrial control systems, medical devices, and legacy IoT deployments [22].

### *B. Cryptographic Agility*

Cryptographic agility, the capacity of a system to switch between cryptographic algorithms without requiring fundamental architectural redesign, has emerged as a central design principle for post-quantum migration, distinct from the question of which specific algorithm to adopt. Systems that hard-code assumptions about key sizes, signature lengths, or specific algorithm identifiers into protocol specifications, application logic, or hardware implementations face substantially higher migration costs than systems designed with abstraction layers that treat the underlying cryptographic primitive as a replaceable component. The selection of multiple structurally independent algorithms by NIST, lattice-based ML-KEM and ML-DSA alongside hash-based SLH-DSA and prospective code-based HQC, reflects an institutional recognition that agility extends beyond protocol design to algorithm selection itself: an organization that can rapidly pivot from one mathematical foundation to another in response to a future cryptanalytic breakthrough is better positioned than one committed irreversibly to a single approach.

### *C. Legacy Systems and IoT Constraints*

The Internet of Things presents a particularly acute version of the migration challenge because many deployed devices were designed with hardware resource budgets calculated against classical cryptographic requirements, with little or no margin for the increased computational, memory, and energy demands of post-quantum algorithms, and because many such devices have operational lifespans measured in decades with limited or no capacity for firmware updates after deployment [22], [23]. Smart cards, industrial control systems, and embedded medical devices share this constraint profile. For newly designed systems, the available remedy is straightforward in principle, select appropriate lightweight parameter sets and build in update capability, even if execution remains an engineering and cost challenge. For already-deployed legacy systems lacking remote update capability, the only remaining options are often physical replacement, which carries direct cost, or acceptance of residual quantum risk for the remaining operational life of the device, a risk-acceptance decision that should be made deliberately rather than by default.

### *D. Organizational and Economic Barriers*

Beyond the technical dimensions of migration, organizations face a coordination problem that is at least as significant as the engineering challenge. Cryptographic dependencies are frequently undocumented, embedded in third-party libraries, vendor firmware, and legacy protocols whose original implementers may no longer be reachable, making the initial inventory step, identifying where and how cryptography is used across an organization's systems, a substantial undertaking in its own right before any migration work can begin [24]. Supply chain coordination compounds this difficulty, since an organization's own migration may be gated on upstream vendors and downstream customers completing compatible upgrades, a dependency structure historically associated with protracted, multi-year transitions in other large-scale cryptographic migrations, such as the industry-wide move away from SHA-1 [24]. The economic calculus of migration timing is therefore not purely technical: organizations must weigh the cost of early action, including potential rework if standards continue to evolve through the ongoing fourth-round NIST process, against the cost of delay, including exposure to harvest-now-decrypt-later collection and the risk of being forced into reactive, more expensive migration under future regulatory deadlines.

### *E. Sector-Specific Migration Considerations*

The practical shape of migration differs considerably across sectors, reflecting differences in data sensitivity horizons, regulatory exposure, and the prevalence of legacy or embedded systems. In the financial sector, public-key cryptography underlies not only transaction authentication but also blockchain-based systems whose elliptic curve foundations make them a direct target of the ECC-focused algorithmic advances described in Section III-B; the prospect, however distant, of a future quantum attack against secp256k1-class curves has already prompted preliminary

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1288

discussion within parts of the digital asset industry about post-quantum signature schemes for next-generation protocols. In healthcare, genomic and clinical data carry confidentiality requirements that can extend for the lifetime of a patient and, in the case of heritable genomic information, for the lifetimes of descendants, placing this sector among those most exposed under the harvest-now-decrypt-later model even though healthcare information systems are not typically considered high-value targets in conventional cybersecurity threat modeling. In government and defense, the combination of long classification periods and the high strategic value of intercepted communications has made this sector an early mover in post-quantum planning, consistent with the G7's critical-infrastructure-first prioritization described in Section VI-B. In industrial and critical-infrastructure contexts, including energy and water systems, the dominant constraint is less data sensitivity than the embedded-systems performance overhead documented in Section V-A, since control-system hardware is frequently resource-constrained and deployed with multi-decade operational lifespans that predate any consideration of post-quantum requirements at design time.

## VI. POLICY AND GOVERNANCE RESPONSES

### *A. NIST Migration Guidance (IR 8547)*

NIST published an initial public draft of Special Publication IR 8547, "Transition to Post-Quantum Cryptography Standards," in November 2024, providing structured guidance on migration sequencing and explicitly cataloguing which existing cryptographic standards and protocols require replacement under the new PQC framework [25]. Notably, this guidance clarifies that symmetric cryptographic standards, including AES, are not expected to require replacement as part of the PQC migration, since their primary quantum-era exposure, the Grover's algorithm threat described in Section II-C, is adequately addressed through existing key-length provisions rather than requiring new algorithms [25]. This distinction is operationally important because it allows organizations to scope migration efforts specifically to public-key infrastructure, key exchange protocols, and digital signature systems, rather than treating the transition as a wholesale replacement of all cryptography in use.

### *B. G7 and International Coordination*

In January 2026, the G7 Cyber Expert Group issued a coordination statement identifying 2035 as an overall target date for completing the transition to quantum-resistant cryptography across G7 jurisdictions, while specifically prioritizing critical systems, financial infrastructure, government communications, and essential services, for completion within a tighter 2030 to 2032 window [26]. This two-tier target structure reflects an explicit policy judgment that not all systems carry equal urgency: critical infrastructure, where the consequences of a future quantum-enabled compromise would be most severe or where data confidentiality requirements extend furthest into the future, is prioritized ahead of lower-risk systems for which a longer migration runway is judged acceptable. The emergence of explicit multilateral coordination on migration timing, rather than each jurisdiction setting independent deadlines, suggests an attempt to avoid the kind of fragmented, inconsistent transition that complicated earlier cryptographic migrations conducted without comparable international alignment.

### *C. Sector-Specific Regulatory Movement*

Beyond the general G7 coordination statement, sector-specific regulators in finance, healthcare, and critical infrastructure have begun incorporating post-quantum readiness into existing cybersecurity compliance frameworks, reflecting recognition that sectors handling long-lived sensitive data, financial transaction histories, genomic and health records, and infrastructure design specifications, face disproportionate harvest-now-decrypt-later exposure relative to sectors whose data has shorter confidentiality shelf lives. The practical effect of this regulatory movement is to convert post-quantum migration from a purely voluntary best practice into an emerging compliance obligation, a shift that is likely to accelerate organizational investment in migration planning over the remainder of the decade.

## VII. DISCUSSION

### *A. Risk Assessment Synthesis*

The evidence reviewed in this paper supports a risk assessment with two seemingly opposed but in fact complementary conclusions. On one hand, no cryptographically relevant quantum computer exists today, and the gap between current hardware, fielding at most a few dozen logical qubits, and the thousands of logical qubits required to threaten RSA-2048 or comparable ECC key sizes remains substantial, measured in hardware generations rather than incremental improvements [9]. Claims of imminent, near-term capability, such as the disputed March 2026 JVG algorithm announcement, have not withstood scientific scrutiny and should be weighted accordingly [13]. On the other hand, the consistent and substantial downward revision of resource estimates achieved through verified algorithmic improvements, the roughly twenty-fold reduction in physical qubits needed for ECC and the broader reduction from twenty million to under one million physical qubits for RSA-2048, demonstrates that the threat timeline is more sensitive to intellectual progress than to brute hardware scaling, and intellectual progress is inherently less predictable and potentially faster than engineering scale-up [1], [2], [7]. These two facts together justify the policy consensus reflected in NIST and G7 guidance: treat migration as urgent and present-tense for high-value, long-lived data, without treating speculative capability claims as confirmed threats.

### *B. Recommendations for Practitioners*

Organizations beginning post-quantum migration planning should prioritize, in sequence, a comprehensive cryptographic inventory to establish where and how public-key cryptography is used across systems and third-party dependencies; a data-sensitivity classification exercise that identifies which datasets carry confidentiality requirements extending far enough into the future to be exposed under the harvest-now-decrypt-later model; risk-prioritized migration of the highest-exposure systems first, consistent with the

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1288

tiered approach reflected in the G7's critical-infrastructure-first timeline; and architectural investment in cryptographic agility so that future algorithm updates, including the pending FN-DSA standard and fourth-round KEM alternatives such as HQC, can be incorporated without renewed wholesale system redesign. For resource-constrained environments, including IoT and embedded systems, organizations should conduct deployment-specific performance and energy benchmarking before parameter selection, given the multiple-order-of-magnitude performance variation documented across hardware classes in Section V-A, rather than assuming a single PQC configuration will be uniformly viable.

*C. Limitations*

This paper's assessment is constrained by the inherent uncertainty of forecasting a rapidly evolving technical field; both quantum hardware capability and the algorithmic landscape for cryptanalysis are subject to revision, and any synthesis of the threat timeline reflects a snapshot of expert opinion and published results as of mid-2026 rather than a stable ground truth. The disputed status of the JVG algorithm claim illustrates a broader methodological challenge for threat assessment in this domain: distinguishing verified results from premature or unsubstantiated announcements requires ongoing engagement with primary technical literature and the scientific response to it, a process that is necessarily provisional and subject to revision as peer review proceeds. Additionally, this paper's treatment of migration economics is qualitative; a fuller quantitative cost-benefit analysis of migration timing under varying threat-timeline assumptions would be a valuable direction for future research, as would empirical study of cryptographic agility implementations across different protocol and system architectures.

A further limitation concerns scope: this paper focuses on the threat to public-key cryptography and the NIST standardization response, and does not provide comprehensive treatment of quantum key distribution (QKD), a physically distinct approach to quantum-resistant communication that uses quantum mechanical properties to detect eavesdropping rather than relying on computational hardness assumptions. QKD raises a different set of deployment considerations, including specialized hardware requirements and distance limitations, that fall outside the scope of the standards-based migration pathway examined here but represent a complementary area of active research and limited real-world deployment, principally in specialized high-security network segments.

*D. Future Research Directions*

Several directions emerge from the gaps identified above as priorities for continued research. First, empirical longitudinal studies tracking organizational post-quantum migration progress against the G7 and sector-specific timelines described in Section VI would provide an evidentiary basis for assessing whether current policy targets are achievable, an assessment that is currently more aspirational than empirically grounded. Second, further work on lightweight and lattice-based cryptographic implementations specifically optimized for the most constrained classes of embedded hardware would help close the three-order-of-magnitude performance gap documented in Section V-A, which currently represents one of the most concrete near-term obstacles to comprehensive migration. Third, continued research into cryptographic diversity, including the maturation of code-based alternatives such as HQC, remains important as a hedge against the possibility, currently considered unlikely but not excluded, that future cryptanalysis identifies an unforeseen weakness in lattice-based hardness assumptions. Fourth, and reflecting the central methodological theme of this paper, the research community would benefit from more formalized mechanisms for rapidly and transparently adjudicating contested capability claims of the kind exemplified by the JVG episode, given the demonstrated speed at which such claims can influence public perception and, potentially, policy and investment decisions before scientific review can run its course.

## VIII. CONCLUSION

The threat that quantum computing poses to classical public-key cryptography has moved, over the period covered by this paper, from a distant theoretical concern to an active and internationally coordinated policy priority, even though no cryptographically relevant quantum computer yet exists. This apparent tension is resolved by recognizing that the urgency of migration derives less from precise timeline forecasting than from the structural features of the threat: the harvest-now-decrypt-later model exposes long-lived data to future decryption regardless of exactly when that capability arrives, and the consistent, algorithm-driven compression of resource estimates over the past two years demonstrates that the relevant timeline is more volatile than hardware-centric forecasts alone would suggest. NIST's finalized post-quantum cryptography standards, FIPS 203, 204, and 205, provide the technical foundation for migration, while NIST's own transition guidance and the G7's January 2026 coordination statement provide an emerging policy framework for sequencing that migration by risk and criticality. The principal remaining challenges are not cryptographic but organizational and economic: building accurate cryptographic inventories, designing for algorithmic agility, and managing the substantial performance overheads documented on constrained hardware. Organizations that treat post-quantum migration as a present-tense operational program, prioritized by data sensitivity and system criticality rather than deferred pending timeline certainty, will be best positioned to manage a threat whose exact arrival date remains genuinely unknown but whose eventual realization is treated by the overwhelming weight of expert opinion as a matter of when, not if.

## REFERENCES

[1]. C. Gidney, "How to factor 2048 bit RSA integers with fewer than one million noisy qubits," arXiv preprint, 2025.

[2]. Google Quantum AI, "Reducing the cost of breaking elliptic-curve cryptography with quantum computers," Google Research Blog, Mar. 31, 2026.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1288

[3]. Caltech-Berkeley-Oratomic Collaboration, "Neutral-atom architectures for resource-efficient implementation of Shor's algorithm," arXiv preprint, Mar. 2026.
[4]. National Institute of Standards and Technology, "NIST releases first 3 finalized post-quantum encryption standards," NIST News, Aug. 13, 2024.
[5]. Authors, "Securing cryptography in the age of quantum computing and AI: Threats, implementations, and strategic response," arXiv preprint arXiv:2603.06969, 2026.
[6]. Authors, "Towards quantum-resistant trusted computing: Architectures for post-quantum integrity verification techniques," arXiv preprint arXiv:2601.11095, 2026.
[7]. C. Gidney and M. Ekera, "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits," Quantum, vol. 5, p. 433, 2021.
[8]. National Institute of Standards and Technology, "NIST IR 8547 (Initial Public Draft): Transition to post-quantum cryptography standards," Nov. 2024.
[9]. Quantum Zeitgeist, "Cryptographically relevant quantum computer: Complete 2026 guide," 2026.
[10]. ScienceAlert, "Quantum computers could break encryption far sooner than we realized," Apr. 13, 2026.
[11]. The Conversation, "Quantum computers are coming to break our codes faster than anyone expected," Apr. 12, 2026.
[12]. Advanced Quantum Technologies Institute, "AQTI announces JVG algorithm for efficient integer factorization," Press Release via PR Newswire, Mar. 2, 2026.
[13]. Diplotic, "Fact check: Are quantum computers now breaking RSA encryption?," Mar. 17, 2026.
[14]. Global Risk Institute, "Quantum threat timeline report," Annual Expert Survey, 2026.
[15]. A. Gautam, "Quantum computers just got 1,000x more efficient at breaking RSA: Developers have less time than they think," Mar. 5, 2026.
[16]. CNN, "Quantum computing threatens to unleash a cybersecurity crisis," May 17, 2026.
[17]. IBM Newsroom, "IBM-developed algorithms announced as NIST's first published post-quantum cryptography standards," PR Newswire, Aug. 13, 2024.
[18]. National Institute of Standards and Technology, "Module-Lattice-Based Key-Encapsulation Mechanism Standard," Federal Information Processing Standards Publication 203, Aug. 2024.
[19]. Sectigo, "NIST's official 2024 post-quantum algorithms," Sectigo Resources, 2025.
[20]. Authors, "Towards quantum-resistant trusted computing," arXiv preprint arXiv:2601.11095, 2026.
[21]. National Institute of Standards and Technology, "NIST IR 8547 (ipd): Transition to post-quantum cryptography standards," Sec. 4.2, Nov. 2024.
[22]. Authors, "Securing cryptography in the age of quantum computing and AI," arXiv preprint arXiv:2603.06969, 2026.
[23]. Palo Alto Networks, "What are NIST PQC standards?," Cyberpedia, 2025.
[24]. National Institute of Standards and Technology, "NIST IR 8547 (ipd): Transition to post-quantum cryptography standards," migration sequencing section, Nov. 2024.
[25]. National Institute of Standards and Technology, "NIST IR 8547 (Initial Public Draft): Transition to post-quantum cryptography standards," Nov. 2024.
[26]. G7 Cyber Expert Group, "G7 coordination statement on post-quantum cryptographic migration," Jan. 2026.