

Quantum Computing and Artificial Intelligence Security

Rizwan Tanveer

PhD Researcher, Doctor of Philosophy in Technology (Cybersecurity)

Walsh College, Troy, Michigan, United States

Principal Consultant, Cybersecurity GRC, Data Privacy and AI Governance

DTS Solution (Beyon Cyber), Dubai, United Arab Emirates

Email: rtanveer@mail.walshcollege.edu

ORCID: 0009-0001-0984-7764

Version 1.0, May 2026

Abstract

Background: The maturation of post-quantum cryptography and the rapid deployment of agentic and generative artificial intelligence are converging into a single, under-examined risk surface. The United States National Institute of Standards and Technology finalised its first three post-quantum cryptographic standards in August 2024 and selected a further algorithm for standardisation in March 2025, yet the security discourse continues to treat quantum cryptanalysis and AI system security as separate domains.

Purpose: This paper consolidates the quantum dimension of AI security into a coherent analytical frame. It examines three intersecting threat vectors, namely the cryptographic exposure of AI assets to harvest-now-decrypt-later strategies, the prospective acceleration of adversarial machine learning by quantum optimisation, and the bidirectional synergy through which each technology amplifies the offensive potential of the other, and it proposes a migration governance agenda suited to practitioners operating multi-jurisdictional AI estates.

Approach: The study adopts a conceptual and integrative review methodology, synthesising primary standards documentation, national authority guidance, and emerging peer-reviewed and preprint literature. It maps identified threats against established control frameworks, including the NIST AI Risk Management Framework, ISO/IEC 42001, and the cryptographic provisions of allied national guidance, and develops a tiered migration model calibrated to data longevity and asset criticality.

Findings: AI estates present a distinctive and elevated harvest-now-decrypt-later exposure because model weights, proprietary training corpora, and inference traffic carry long confidentiality lifespans that frequently exceed plausible timelines for a cryptographically relevant quantum computer. Cryptographic agility, rather than any single algorithm choice, emerges as the decisive architectural property. The convergence of quantum acceleration and adversarial machine learning is currently theoretical but warrants anticipatory governance, given the asymmetry between preparation costs and tail risk.

Implications: Organisations should treat post-quantum migration as an AI governance obligation rather than a narrow cryptographic upgrade, embedding cryptographic inventory, agility, and data-longevity triage within existing AI management systems. For the Gulf Cooperation Council region, where sovereign data initiatives and AI adoption are advancing in parallel, early alignment with internationally recognised post-quantum standards offers both a risk-reduction and a strategic-positioning advantage.

Keywords: post-quantum cryptography; AI security; harvest-now-decrypt-later; cryptographic agility; adversarial machine learning; quantum risk governance; ML-KEM

JEL Classification: O33; K24; L86

ACM CCS: Security and privacy ~ Public key (asymmetric) techniques; Security and privacy ~ Cryptanalysis and other attacks; Computing methodologies ~ Machine learning

1. Introduction

Two technological trajectories that have, until recently, occupied separate corners of the security literature are now converging. The first is the maturation of quantum computing and the corresponding standardisation of cryptography designed to withstand it. The second is the rapid operationalisation of artificial intelligence, increasingly in agentic and multi-agent configurations that hold, process, and transmit assets of considerable and durable value. This paper argues that the intersection of these trajectories constitutes a distinct risk surface, one that is inadequately addressed when quantum cryptanalysis and AI system security are treated as unrelated concerns.

The standardisation milestone is no longer prospective. The United States National Institute of Standards and Technology (NIST) published its first three finalised post-quantum cryptographic standards on 13 August 2024: FIPS 203, specifying a module-lattice-based key-encapsulation mechanism derived from the CRYSTALS-KYBER submission and now designated ML-KEM; FIPS 204, specifying a module-lattice-based digital signature standard (ML-DSA) derived from CRYSTALS-Dilithium; and FIPS 205, a stateless hash-based signature standard (SLH-DSA) derived from SPHINCS+. A fourth algorithm, HQC, was subsequently selected for standardisation in March 2025 to provide a key-establishment mechanism resting on a mathematical foundation distinct from the lattice family, thereby reducing systemic reliance on a single class of hard problem. The selection phase is effectively concluded; the migration phase has begun.

Yet the prevailing framing of post-quantum migration as a cryptographic refresh, carried out by network and infrastructure teams and largely invisible to AI governance, understates the problem for organisations whose principal assets are AI systems. The confidentiality lifespan of a proprietary model, its training corpus, and its accumulated inference history frequently exceeds the interval before a cryptographically relevant quantum computer is widely expected to emerge. This temporal mismatch is the crux of the harvest-now-decrypt-later threat, and it bears with particular force on AI estates.

This paper makes three contributions. First, it consolidates the quantum dimension of AI security into a single analytical frame, distinguishing three intersecting threat vectors that the literature has tended to treat in isolation. Second, it maps these vectors onto established AI governance instruments, demonstrating that post-quantum readiness can and should be expressed as an obligation within an AI management system rather than as a parallel cryptographic project. Third,

it advances a tiered migration governance agenda calibrated to data longevity and asset criticality, with specific attention to the Gulf Cooperation Council (GCC) operating context. The paper situates itself within a wider series on the security of AI systems and connects directly to earlier treatments of adversarial machine learning and supply chain security.

2. The Quantum Threat to Classical Cryptography

2.1 Shor, Grover, and the asymmetry of exposure

The quantum threat to contemporary cryptography is not uniform across primitives, and conflating its effects obscures sound prioritisation. Shor's algorithm, when run on a sufficiently large and fault-tolerant quantum computer, solves the integer factorisation and discrete logarithm problems in polynomial time, rendering RSA, finite-field Diffie-Hellman, and elliptic-curve cryptography effectively broken. This is the existential threat, and it falls on the asymmetric primitives that underpin key establishment, digital signatures, and the public-key infrastructure on which almost all secure communication depends.

Grover's algorithm presents a materially different and more contained problem. It offers a quadratic speed-up against unstructured search, which halves the effective security level of symmetric ciphers and hash functions. The practical remedy is correspondingly modest: doubling key and digest lengths, for instance moving from AES-128 to AES-256, restores the original security margin. The widely repeated claim that quantum computing breaks all cryptography is therefore inaccurate and strategically unhelpful. The urgent exposure is concentrated in asymmetric cryptography, and migration effort should be prioritised accordingly.

2.2 The standardised post-quantum portfolio

The finalised NIST standards provide a deployable response to the asymmetric exposure. The portfolio is deliberately diversified across functions and mathematical bases, as summarised below. The diversification is itself a security property: a cryptanalytic advance against the lattice family would not simultaneously compromise the hash-based signature standard or the code-based key-establishment mechanism.

Standard	Designation	Function	Mathematical basis
FIPS 203	ML-KEM	Key encapsulation	Module lattices (MLWE)
FIPS 204	ML-DSA	Digital signature (primary)	Module lattices (MLWE)

Standard	Designation	Function	Mathematical basis
FIPS 205	SLH-DSA	Digital signature (stateless)	Hash functions
HQC (selected 2025)	—	Key encapsulation (backup)	Error-correcting codes

Table 1. The standardised post-quantum portfolio and its deliberate diversification across function and mathematical foundation.

National authority guidance has begun to translate these standards into migration timelines. The United States National Security Agency's Commercial National Security Algorithm Suite 2.0 mandates the use of a lattice-based key-encapsulation mechanism for national security systems, with phased timelines extending into the early 2030s for legacy infrastructure. Allied guidance, including that issued by the United Kingdom's National Cyber Security Centre, endorses the NIST standards for interoperability and recommends minimum security categories for commercial use. The direction of travel is unambiguous, and procurement-driven compliance pressure is already propagating through supply chains.

3. Why AI Estates Face Distinctive Quantum Exposure

3.1 The harvest-now-decrypt-later threat model

The harvest-now-decrypt-later strategy, sometimes termed store-now-decrypt-later, requires no current quantum capability whatsoever. An adversary needs only to intercept and retain encrypted material today, in the expectation that a cryptographically relevant quantum computer will later render it readable. The attack, therefore, decouples the moment of compromise from the moment of capability: confidentiality is breached retroactively, and the breach is invisible at the time the data is harvested. Joint guidance from national authorities has warned that well-resourced actors, particularly nation-states, are plausibly already conducting such collection at scale.

The decisive variable in this model is the confidentiality lifespan of the harvested data relative to the expected arrival of quantum cryptanalysis. Data whose value expires within a year or two carries limited exposure. Data that must remain confidential for a decade or more is already at risk, because it may still be sensitive when decryption becomes feasible. This is conventionally expressed through Mosca's inequality, which holds that an organisation should act when the sum of the time its data must remain secure, and the time required to migrate its systems, exceeds the time until a quantum computer can break the prevailing cryptography.

3.2 The elevated exposure of AI assets

AI estates are distinctive precisely because their core assets combine high value with long confidentiality lifespans. Three categories warrant particular attention:

- **Model weights.** A trained foundation or domain-specific model represents a substantial and durable capital investment. Its parameters retain commercial and, in some settings, national-security value for many years. Weights exfiltrated in encrypted form today and later decrypted would expose proprietary capabilities long after the harvesting event, with no contemporaneous indication of compromise.
- **Training corpora.** Proprietary and regulated training data, which may include personal, health, or sovereign records, frequently carries statutory confidentiality obligations measured in decades. Harvested training sets thus represent both a competitive and a compliance liability with an unusually long tail.
- **Inference traffic and accumulated interaction history.** The prompts, responses, and tool-invocation records exchanged with deployed AI systems can encode highly sensitive operational, commercial, and personal information. In agentic configurations, this traffic also reveals the structure of an organisation's automated decision-making, which has intelligence value independent of the data content.

The implication is that AI governance cannot treat post-quantum migration as someone else's infrastructure problem. The assets most exposed to harvest-now-decrypt-later are frequently the very assets an AI management system exists to protect. This reframing, from cryptographic housekeeping to AI governance obligation, is the central argument of this paper.

4. Adversarial Convergence: Quantum and AI as Mutual Amplifiers

Beyond the cryptographic exposure of static assets, a second and more speculative intersection concerns the dynamic amplification of attacks. The relationship between quantum computing and AI is bidirectional, and each direction carries security consequences that anticipatory governance should acknowledge, even where near-term feasibility remains uncertain.

4.1 Quantum acceleration of adversarial machine learning

Adversarial machine learning is an established weakness of contemporary models, in which inputs perturbed imperceptibly to a human observer cause confident misclassification. The generation of

such adversarial examples is fundamentally an optimisation problem, namely the search for a minimal perturbation that crosses a decision boundary. Where optimisation is the bottleneck, quantum approaches that accelerate search and sampling could, in principle, compress the cost of crafting effective adversarial inputs and broaden the feasible attack surface against deployed models. The claim should be stated with care: this is a prospective amplification of a known attack class, not a demonstrated capability at scale, and current quantum hardware remains noisy and limited. It nonetheless merits inclusion in a forward-looking threat model precisely because the underlying adversarial vulnerability is already real and widespread.

4.2 AI acceleration of cryptanalysis and quantum engineering

The converse direction is equally consequential. Machine learning is increasingly applied to improve quantum systems themselves, optimising qubit control, assisting quantum error correction, and supporting algorithm design. AI-driven cryptanalysis, while not a substitute for a fault-tolerant quantum computer, may also sharpen side-channel analysis and accelerate the discovery of implementation weaknesses in deployed cryptography, including in early and imperfect post-quantum deployments. The migration to post-quantum standards does not, therefore, close the cryptanalytic frontier; it relocates it, and implementation quality becomes a first-order security concern.

4.3 A note on proportionality and hype

A balanced treatment must resist both complacency and alarmism. Vendor discourse frequently overstates the imminence of both quantum cryptanalysis and autonomous AI attack capability. The proportionate response rests on an asymmetry argument rather than on a confident timeline: the cost of cryptographic agility and data-longevity triage is modest and largely consists of good architectural hygiene, whereas the tail risk of unpreparedness, in the specific case of long-lived AI assets, is severe and retroactive. Action is justified by the structure of the risk, not by any prediction of when a cryptographically relevant quantum computer will arrive.

5. Mapping Quantum Readiness onto AI Governance Frameworks

The argument that post-quantum migration is an AI governance obligation is operationally meaningful only if it can be expressed through the instruments practitioners already use. This

section maps quantum-readiness controls onto established frameworks, demonstrating that no parallel governance structure is required.

Framework	Relevant locus	Quantum-readiness expression
NIST AI RMF	Map and Manage functions; data and asset inventory	Cryptographic inventory of AI assets; data-longevity triage as a risk-mapping activity
ISO/IEC 42001	AI management system controls, lifecycle, and supplier management	Cryptographic agility as a documented control objective; supplier post-quantum attestation
NIST AI 600-1 (GenAI profile)	Information security and confidentiality risks	HNDL exposure of weights, corpora, and inference traffic treated as a named risk
EU AI Act	Cybersecurity obligations for high-risk and general-purpose models	State-of-the-art cryptography interpreted to include post-quantum readiness over the system's lifespan

Table 2. Quantum-readiness controls expressed within established AI governance frameworks, requiring no parallel governance structure.

The mapping yields a single recurring control that subsumes the others: cryptographic agility, understood as the architectural capacity to substitute cryptographic algorithms without redesigning the systems that depend on them. Because the post-quantum standards are themselves diversified and may be supplemented further, and because hybrid classical-and-post-quantum deployments are the prudent transitional posture, agility rather than any specific algorithm choice is the decisive and durable property. An AI management system that institutionalises cryptographic inventory, agility, and data-longevity triage is, by construction, quantum-ready, and remains so as the standards evolve.

6. A Tiered Migration Governance Agenda

The following agenda translates the preceding analysis into a sequenced, risk-prioritised programme. It is deliberately calibrated to data longevity and asset criticality rather than to a uniform deadline, reflecting the logic of Mosca's inequality and the proportionality argument advanced above.

6.1 Tier 1: Inventory and triage (immediate)

The foundational step is a cryptographic inventory of the AI estate that records where classical asymmetric cryptography protects model weights, training corpora, inference channels, and supporting infrastructure, and that classifies each asset by required confidentiality lifespan. Long-lived, high-criticality assets, typically regulated training data and high-value model weights, are flagged for priority migration. This activity fits naturally within the mapping function of the NIST AI RMF and the asset management controls of ISO/IEC 42001.

6.2 Tier 2: Agility and hybrid deployment (near term)

Systems should be re-architected for cryptographic agility, decoupling algorithm choice from application logic. Where confidentiality lifespans justify it, hybrid key-establishment that combines a classical and a post-quantum mechanism should be deployed, providing protection that holds if either component remains unbroken. Priority is given to data at rest for long-lived assets and to data in transit for sensitive inference channels, since these are the loci of greatest harvest-now-decrypt-later exposure.

6.3 Tier 3: Supply chain and attestation (near to medium term)

Post-quantum readiness must extend across the AI supply chain, encompassing foundation-model providers, orchestration platforms, and tool integrators. Procurement and supplier-management processes should require post-quantum migration roadmaps and, over time, attestation of quantum-safe key establishment and signing. This tier connects directly to supply chain security treatments elsewhere in this series and to the orchestration-liability question of who, across a layered AI stack, bears responsibility for cryptographic posture.

6.4 Tier 4: Monitoring and re-evaluation (continuous)

Because both the quantum timeline and the cryptanalytic standing of individual post-quantum algorithms remain subject to revision, the programme must include continuous monitoring of standards developments, national-authority guidance, and any cryptanalytic results affecting deployed primitives. The diversified standard portfolio and an agile architecture together ensure that a future compromise of one algorithm family can be absorbed through substitution rather than redesign.

7. The Gulf Cooperation Council Operating Context

The GCC region presents a distinctive vantage on quantum-AI security risk. Sovereign data initiatives, national AI strategies, and rapid public-sector and critical-infrastructure adoption of AI are advancing in parallel, often within jurisdictions that have recently enacted or are developing data-protection and AI-governance regimes. Two features sharpen the relevance of the migration agenda in this setting.

First, sovereign and regulated data holdings, including population-scale health, financial, and government records that increasingly feed AI systems, carry long statutory confidentiality lifespans and are therefore prime targets for harvest-now-decrypt-later. The temporal mismatch at the heart of the threat model is especially acute where data sovereignty is a stated policy objective. Second, the region's posture as a fast adopter of advanced AI creates an opportunity to embed post-quantum readiness at the point of system design rather than retrofitting it, which is both cheaper and more robust. Early alignment with internationally recognised post-quantum standards therefore offers a combined risk-reduction and strategic-positioning advantage, supporting interoperability with allied systems while reinforcing sovereign-data commitments.

For practitioners advising GCC clients, the practical recommendation is to fold post-quantum readiness into AI governance engagements from the outset, treating cryptographic inventory and agility as standard components of an AI management system rather than as a separate and deferrable cryptographic project.

8. Conclusion

The convergence of post-quantum cryptography and AI security is no longer a horizon concern. With the foundational standards finalised and migration timelines propagating through national guidance and procurement, the open question is not whether to act but how to prioritise. This paper has argued that AI estates face a distinctive and elevated harvest-now-decrypt-later exposure, because their core assets, model weights, training corpora, and inference traffic, combine high value with long confidentiality lifespans that frequently outrun plausible quantum timelines.

The principal contribution is a reframing: post-quantum migration is an AI governance obligation, expressible through the NIST AI RMF, ISO/IEC 42001, and allied instruments, rather than a parallel cryptographic project conducted out of sight of those accountable for AI risk.

Cryptographic agility, not any single algorithm, is the decisive and durable architectural property, and a tiered migration agenda calibrated to data longevity and asset criticality provides a proportionate path forward. The adversarial convergence of quantum acceleration and AI attack techniques remains prospective, but the asymmetry between modest preparation costs and severe retroactive risks justifies anticipatory governance now.

Subsequent work in this series turns to cross-jurisdictional compliance and non-human identity and access management, both of which intersect with the quantum-readiness agenda developed here: the former through divergent national timelines that complicate multi-jurisdictional migration, and the latter through the post-quantum re-establishment of machine and agent credentials on which agentic AI security ultimately depends.

Declarations

Competing interests. The author is Principal Consultant in Cybersecurity GRC, Data Privacy and AI Governance at DTS Solution (Beyon Cyber), an organisation that provides cybersecurity advisory services in the Gulf Cooperation Council region. The author has no financial or other personal interest in any product, vendor, AI system, or regulatory framework discussed in this paper.

Funding. This research received no external funding. The work was conducted as part of the author's doctoral research at Walsh College.

Data availability. Not applicable. This paper is a narrative review and conceptual synthesis; no primary data were collected.

Ethics approval. Not applicable; the study involved no human participants.

References

- Cloud Security Alliance. (2024, August 15). NIST FIPS 203, 204 and 205 finalized: An important step towards a quantum-safe future.
<https://cloudsecurityalliance.org/blog/2024/08/15/nist-fips-203-204-and-205-finalized-an-important-step-towards-a-quantum-safe-future>
- Greshake, K., Abdelnabi, S., Mishra, S., Endres, C., Holz, T., & Fritz, M. (2023). Not what you've signed up for: Compromising real-world LLM-integrated applications with indirect prompt injection. Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security.
- National Institute of Standards and Technology. (2024). FIPS 203: Module-lattice-based key-encapsulation mechanism standard. U.S. Department of Commerce.
<https://doi.org/10.6028/NIST.FIPS.203>
- National Institute of Standards and Technology. (2024). FIPS 204: Module-lattice-based digital signature standard. U.S. Department of Commerce.
<https://doi.org/10.6028/NIST.FIPS.204>
- National Institute of Standards and Technology. (2024). FIPS 205: Stateless hash-based digital signature standard. U.S. Department of Commerce.
<https://doi.org/10.6028/NIST.FIPS.205>
- National Institute of Standards and Technology. (2024). Artificial intelligence risk management framework: Generative artificial intelligence profile (NIST AI 600-1). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.600-1>
- National Institute of Standards and Technology. (2025). Status report on the fourth round of the NIST post-quantum cryptography standardization process (NIST IR 8545). U.S. Department of Commerce.
- National Security Agency. (2022). Announcing the Commercial National Security Algorithm Suite 2.0 (CNSA 2.0). [Cybersecurity advisory].
- International Organization for Standardization. (2023). ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system. ISO.

- Mosca, M. (2018). Cybersecurity in an era with quantum computers: Will we be ready? *IEEE Security & Privacy*, 16(5), 38–41. <https://doi.org/10.1109/MSP.2018.3761723>
- Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5), 1484–1509. <https://doi.org/10.1137/S0097539795293172>
- Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, 212–219. <https://doi.org/10.1145/237814.237866>
- Tanveer, R. (2026). Adversarial machine learning: Attack vectors, defences, and robustness [Working paper]. SSRN. [Companion paper, this series.]
- Tanveer, R. (2026). Supply chain security for artificial intelligence systems [Working paper]. SSRN. [Companion paper, this series.]