

The Quantum AI Convergence in Cybersecurity: How Artificial Intelligence Is Accelerating Both the Quantum Threat and the Post-Quantum Defense

Troy Coienth Troublefield^{1,2,3} 

¹Department of Cyberpsychology, Capital Technology University, Laurel, MD, USA

²Department of Information Technology, Capella University, Minneapolis, MN, USA

³Department of International Business, International School of Management, Paris, France

Email: drtroytroublefield@yahoo.com

How to cite this paper: Troublefield, T.C. (2026) The Quantum AI Convergence in Cybersecurity: How Artificial Intelligence Is Accelerating Both the Quantum Threat and the Post-Quantum Defense. *Journal of Quantum Information Science*, 16, 467-493.

<https://doi.org/10.4236/jqis.2026.164017>

Received: September 6, 2026

Accepted: October 13, 2026

Published: October 16, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Scholarship and enterprise practice have historically treated quantum computing and artificial intelligence (AI) as distinct domains of cybersecurity risk, each with its own timeline, mechanisms, and mitigation literature. This article argues that the separation is no longer analytically or operationally tenable. As of 2026, the two technologies have entered a phase of mutual acceleration in which AI compresses the timeline toward a cryptographically relevant quantum computer (CRQC) while serving as the principal mechanism through which organizations can detect, inventory, and defend against quantum-era threats. The article develops this thesis across five analytical movements: the accelerated offensive timeline, in which machine-learning techniques contribute to advances in error correction and algorithm optimization that have driven estimated qubit requirements for breaking RSA-2048 from roughly twenty million toward one million, a reduction attributable to converging algorithmic, hardware, and AI-assisted advances rather than to any single cause; the harvest-now-decrypt-later (HNDL) threat, which converts a prospective future risk into a present data-exposure problem made more urgent by the state-directed investment postures of well-resourced adversaries; the bidirectional defensive opportunity, in which quantum machine learning (QML), a research-stage methodology whose results are exploratory and conditional on task, dataset, and evaluation environment, and AI-driven cryptographic discovery offer a realistic path to enterprise-scale post-quantum migration; the 2026 regulatory inflection, in which finalized NIST standards and binding compliance deadlines have transformed migration from a research question

into an execution problem; and the global investment and geopolitical competition dimension, in which divergent national quantum programs, military-civil fusion architectures, indigenous cryptographic standards, and supply chain concentration risks expand the scope of the migration challenge beyond the compliance-deadline framing that domestic regulatory guidance alone provides. Drawing on recent empirical results in QML-based intrusion detection, documented advances in AI-assisted quantum error correction, the current standards landscape, and comparative analysis of U.S., Chinese, and European quantum investment postures, the article concludes with a five-part strategic framework for security leaders and identifies directions for further research. The central claim is that the defining cybersecurity challenge of the coming decade resides in the interaction of quantum computing and AI, and in the geopolitical competition shaping both, rather than in either technology considered in isolation.

Keywords

Quantum Computing, Artificial Intelligence, Post-Quantum Cryptography, Harvest Now Decrypt Later, Quantum Machine Learning, Cryptographic Agility, CNSA 2.0, Critical Infrastructure

1. Introduction

Two technologies expected to mature on independent timelines have instead become entangled, and the entanglement has arrived faster than either was individually forecast. Quantum computing has long been treated as a slow-moving, physics-bound problem, its cryptographic consequences deferred to a horizon distant enough that most organizations could plan around it rather than for it. Artificial intelligence, by contrast, has been treated as the technology of the present, already embedded in security operations and already reshaping how both defenders and adversaries' work. This article argues that this separation between them holds. Artificial intelligence is not merely adjacent to the quantum threat; it is now one of the principal forces determining how quickly that threat matures, and it simultaneously furnishes the most practical means available for defending against it.

The argument is deliberately structural rather than alarmist. It does not claim that a cryptographically relevant quantum computer is imminent, nor does it attempt to adjudicate among the competing expert timelines that place that event anywhere from the late 2020s to the 2040s. Its claim is narrower and, for planning purposes, more consequential: the variable that has historically justified deferral, namely how far away the threat sits, is itself being compressed by a technology already deployed at enterprise scale. A posture calibrated to a fixed and distant date is not well suited to a threat whose timeline is being actively revised downward. Compounding this, a substantial share of the risk this convergence produces does not wait for that timeline to resolve, since data harvested and archived today

can be decrypted retroactively once the capability arrives, converting a future contingency into a present exposure.

This reframing carries a second, more constructive implication that runs through the remainder of the article. The same coupling that accelerates the offensive threat also accelerates the defense, since artificial intelligence underwrites both the QML methods now being tested against threat detection and the discovery and migration tooling organizations need to execute a transition that is, in most respects, no longer optional. The practical question this article answers is therefore not whether organizations should be concerned about quantum computing, artificial intelligence, or their interaction, but how a security leader should reason about resource allocation and sequencing once the two are recognized as a single, compound problem rather than two separate ones. The remainder of this introduction lays out the background that motivates that reframing and the scope and sourcing of the analysis that follows.

1.1. Background of the Study

For most of the past decade, enterprise security planning has filed quantum computing and artificial intelligence under separate headings. Quantum computing was understood as a long-horizon cryptographic problem, consequential but situated a decade or more in the future. Artificial intelligence was treated as a near-term operational concern, already reshaping security teams' daily work from phishing generation to automated threat detection. Practitioners could reasonably assign the two technologies to different planning cycles, different budgets, and different specialist teams. Standards bodies reinforced the division, addressing post-quantum cryptography and AI governance in separate publications and on separate schedules. Treating the two as independent was defensible so long as their timelines and underlying mechanisms had little bearing on one another. This article argues that, by 2026, this assumption has become a liability that distorts how organizations allocate attention and resources.

The reason is a coupling that security researchers have increasingly documented: artificial intelligence accelerates quantum computing itself [1]. Machine-learning techniques are now being applied directly to the hardest engineering problems that stand between today's noisy, intermediate-scale quantum processors and a CRQC, including quantum error correction, qubit stabilization, and the optimization of quantum algorithms and circuits [2]. Vendors and academic laboratories alike have reported that machine-learning decoders improve the speed and accuracy of error correction, a central obstacle in the field [2]. As a result, the quantum threat timeline is no longer a fixed, external parameter that organizations can plan around at leisure. It has become an endogenous variable, compressed by a general-purpose technology already deployed at scale across other technical domains [1]. That shift in the character of the timeline, from exogenous constant to moving target, is what forces a corresponding shift in how security leaders should reason about preparation.

This coupling operates in both directions, and that bidirectionality forms the article's analytical core. The same convergence that sharpens the offensive threat also supplies the defense, because AI encompasses both QML for threat detection and AI-driven discovery for cryptographic migration [3]. For most organizations, the defensive applications are more immediately actionable than the offensive threat is imminent, since migration tooling can be deployed against present infrastructure while a CRQC remains prospective. The strategic problem is therefore not a choice between defending against quantum computing and harnessing artificial intelligence. It is managing a single, compound risk surface produced by their interaction. Framing the problem this way changes which investments appear urgent and which appear premature.

The article proceeds in seven sections. Section 2 examines the offensive vector, detailing how AI compresses the quantum timeline and how the HNDL dynamic renders the threat a present concern rather than a future one. Section 3 develops the defensive vector, surveying the empirical state of QML in cybersecurity and AI's operational role in migration. Section 4 situates the analysis within the 2026 regulatory inflection point, where finalized standards and binding deadlines have reframed the problem as one of execution. Section 5 offers a five-part strategic framework for security leaders. Section 6 discusses the theoretical and practical implications of the foregoing analysis, addresses anticipated objections, and notes directions for further research. Section 7 concludes. Each section moves from mechanism to implication, so the strategic recommendations in Section 5 follow from the technical and regulatory analysis that precedes them. The argument is framed throughout around the intersection of the two technologies, on the premise that analyses treating either in isolation misjudge both the urgency and the shape of the response.

1.2. Scope and Approach

This article is a conceptual and integrative contribution rather than an empirical study. Its aim is to synthesize developments that are typically examined in isolation, including advances in quantum hardware and error correction, the maturation of QML for security, the HNDL threat, and the post-quantum standards landscape, into a single analytical frame organized around the interaction of quantum computing and artificial intelligence. The supporting evidence is drawn from three bodies of source material. The first is the recent peer-reviewed and preprint literature on QML for cybersecurity [3]-[6]. The second is technical reporting and primary standards documentation on post-quantum cryptography and quantum hardware progress [2] [7]-[10]. The third is industry and governmental analysis of the threat environment as of 2026 [1] [11] [12]. Where empirical figures are cited, such as qubit-requirement estimates, error-correction milestones, or QML benchmark results, they are reported as the relevant literature presents them, with attention to the limitations the same literature acknowledges. The contribution is the synthesis and the strategic framework that follows, not the generation of new

experimental data.

Two scoping decisions warrant note. First, the article brackets the deep mathematics of post-quantum schemes and the physics of quantum error correction, treating both at the level of their strategic implications rather than their technical internals. This decision reflects the intended audience of security leaders, architects, and policymakers, who require the strategic implications more urgently than the underlying formalism. Readers seeking the mathematical specifications of the standardized algorithms can consult the primary Federal Information Processing Standards publications directly [7]-[9]. Second, the article treats artificial intelligence broadly, encompassing both the machine-learning methods applied to quantum development and the AI tooling applied to cybersecurity defense. This breadth is deliberate, because AI's generality as an accelerant gives the convergence its force. Narrower treatments that consider only one application of AI tend to understate the coupling and, therefore, the urgency of a coordinated response.




Figure 1. The quantum-AI convergence architecture.

Figure 1 visually illustrates the bidirectional relationship between artificial intelligence and quantum cybersecurity. Artificial intelligence sits at the center because it both accelerates offensive quantum capabilities and enables defensive post-quantum security measures. The left pathway depicts AI-driven advances in quantum error correction, algorithm optimization, hardware improvement, and the resulting acceleration toward a cryptographically relevant quantum computer. The right pathway highlights AI's defensive applications, including QML, automated cryptographic discovery, enterprise migration planning, and post-quantum resilience. Both pathways converge to demonstrate that cybersecurity readiness depends upon integrating AI into both offensive risk assessment and defensive modernization strategies.

2. The Offensive Vector: Artificial Intelligence as a Quantum Accelerant

The quantum threat to modern cryptography has long been treated as a distant contingency, deferred on the assumption that a CRQC remains decades away. This section challenges that assumption by examining how artificial intelligence accelerates quantum progress, compressing timelines that organizations have historically used to justify inaction. The analysis proceeds in five parts. It first establishes the formal cryptographic stakes posed by algorithms such as Shor's, and the qubit thresholds long associated with breaking Rivest-Shamir-Adleman (RSA) and Elliptic-curve cryptography (ECC). It then traces how AI-driven advances in error correction and decoding have materially reduced those thresholds, shifting expert probability estimates toward an earlier horizon. The section next turns to HNDL, reframing the quantum threat from a prospective event into a present operational hazard, one that AI intensifies by improving adversarial triage of collected data. A discussion of agentic AI follows, arguing that the expanding attack surface created by autonomous systems is inseparable from the cryptographic migration problem. The section closes by addressing the structure of the timeline debate itself, contending that persistent expert disagreement should be read not as grounds for delay but as evidence of an asymmetric and worsening risk profile that favors early preparation over patient deferral.

2.1. The Cryptographic Stakes

The cryptographic threat posed by quantum computing is well established in its formal mechanics. The public-key systems that secure most digital communication, chiefly RSA and ECC, derive their security from mathematical problems such as integer factorization and the discrete algorithm problem that are believed intractable for classical computers at relevant scales. Shor's algorithm, introduced in 1994, demonstrated that a sufficiently powerful quantum computer could solve these problems in polynomial time, rendering the affected cryptography broken rather than merely weakened [13]. The result reframed public-key cryptography as conditionally secure, its safety contingent on the continued unavailability of large-scale quantum hardware. For three decades, the operative qualifier has been sufficiently powerful, and that qualifier has carried enormous practical weight. Today's quantum processors remain too small and too error-prone to threaten production cryptography. That gap between demonstration and cryptographic relevance is what has allowed most organizations to defer action without apparent consequence.

The distance between present capability and cryptographic relevance has conventionally been expressed in terms of physical qubit counts. Conservative estimates long held that breaking 2048-bit RSA would require on the order of twenty million physical qubits, a figure that placed the threat in the 2035 to 2040 range under most hardware roadmaps [2]. This framing encouraged patience, since the required machine appeared to be two decades away. Two developments have un-

settled that posture. Hardware has advanced materially, as illustrated by Google's Willow processor, announced in late 2024, which demonstrated below-threshold quantum error correction across 105 superconducting qubits and validated that logical error rates can be suppressed exponentially as physical qubit counts grow, a foundational requirement for fault-tolerant computation [14]. Error correction has moved from a theoretical concern to the industry's defining engineering priority, with a marked increase in published research and commercial investment [2]. More consequentially for the timeline, algorithmic and systems-level work has sharply reduced the resources the attack is thought to require, a point developed in the next subsection.

2.2. How AI Compresses the Timeline

The reduction in estimated attack cost is where artificial intelligence enters the analysis. When the central obstacle to a CRQC is a set of research and engineering problems, including the optimization of error-correcting codes, the design of higher-fidelity gates, the refinement of algorithmic implementations, and the real-time decoding of error syndromes, progress is no longer bounded solely by the pace of physics laboratories [2]. Instead, it is augmented by the same machine-learning capabilities advancing across other scientific fields. Researchers already use machine-learning systems to optimize quantum error correction, widely regarded as the field's defining engineering bottleneck [1] [2]. AI-powered decoders in particular have become a recurring feature of the error-correction advances reported through 2025 [2]. Because these decoders address the bottleneck directly rather than incidentally, improvements translate into faster progress toward fault tolerance. The mechanism by which AI accelerates quantum computing is therefore concrete and specific, not a general appeal to technological momentum. Isolating AI's specific contribution from concurrent hardware engineering advances is difficult, and the literature does not yet provide a controlled estimate of how much of the acceleration is attributable to machine-learning decoders versus improvements in fabrication, qubit fidelity, and circuit design. The claim is not that AI is the sole driver, but that it is a demonstrably active contributor to the field's defining bottleneck, error correction, in a way that is identifiable and distinct from hardware scaling alone [1] [2].

The cumulative effect on the threat horizon warrants careful framing. Approaches combining improved software, more efficient algorithms, and higher-quality qubits have driven estimates for breaking RSA-2048 downward, with recent analysis suggesting the task may require about one million physical qubits rather than the previously cited twenty million, a reduction of roughly 95 percent in the estimated resource requirement [2]. Critically, the literature attributes this reduction primarily to algorithmic and systems-level improvements, including more efficient fault-tolerant circuit decompositions and advances in error-correcting codes, rather than to AI in isolation [2]. AI-powered decoders have contributed to faster and more accurate syndrome decoding, one of the field's rate-

limiting engineering tasks, but the full reduction in qubit estimates reflects cumulative advances across multiple research communities. Expert probability estimates for the arrival of a CRQC have shifted correspondingly, with some assessments noting that these estimates have roughly doubled since 2022 [1]. None of this implies that a cryptanalytically relevant machine exists today, and it does not. The planning-relevant claim is narrower: the variable that determines when such a machine arrives is subject to acceleration from multiple converging sources, AI among them, and successive revisions have consistently moved toward sooner rather than later [1] [2]. A single revision of this kind might be dismissed as noise, but a sustained pattern of downward revision in resource estimates and upward revision in probability constitutes a trend. For a planning function, a timeline uncertain in its urgency warrants a different response than one that is simply distant.

2.3. Harvest Now, Decrypt Later

The most consequential near-term expression of the quantum threat requires no functioning quantum computer. HNDL, also termed store now, decrypt later, describes a strategy in which adversaries intercept and archive encrypted data today, holding it in reserve until a CRQC can decrypt it retroactively [12]. Western intelligence services and national cybersecurity authorities have documented this pattern, describing systematic interception and archiving of encrypted traffic by well-resourced actors [12] [15].

The CISA advisory is the most specific primary-source reference available in the public domain on this practice; it does not name actors, quantify collected traffic, or detail operational methods, and broader claims about the systematic scale of such programs rest on analytical inference rather than declassified operational reporting [15]. The HNDL threat is credible and national security authorities treat it as serious, but practitioners should note that the public evidence base describes the plausibility and known incentive structure for such collection rather than providing documented operational specifics [15].

For data whose confidentiality must persist for years, including intellectual property, classified material, health records, financial information, and the training data and model weights of AI systems, the breach has, in an important sense, already occurred [12]. The affected organization may not know it yet and may not learn of it for years. This latency between compromise and disclosure makes HNDL difficult to reason about using conventional incident-response frameworks, which assume a breach becomes known within a bounded period. The threat is not hypothetical harm at some future date but present collection whose consequences are merely deferred.

HNDL reframes quantum risk from a prospective event into a present operational hazard, inverting the logic of patient deferral. The economics of data theft have shifted because a rational adversary is now incentivized to invest in the mass collection of encrypted traffic, expecting decryption capability to arrive within the useful lifetime of that data [12]. Artificial intelligence sharpens the threat along

two additional dimensions. First, it improves an adversary's capacity to triage and prioritize harvested data at scale, identifying which encrypted archives warrant retention and early decryption, which is a non-trivial Problem given the volumes involved.

First, AI may improve an adversary's capacity to triage and prioritize harvested data at scale by enabling automated classification of intercepted data by estimated value, source organization, or content type based on metadata and traffic patterns. Direct evidence that well-resourced state actors are currently deploying machine-learning triage against harvested encrypted archives is not available in the open literature; the claim is a logical inference from the known capabilities of large-scale AI systems applied to data classification problems of comparable scale and structure [12]. The security authorities cited treat the inference as analytically credible, but practitioners should weigh it accordingly [12]. Second, by accelerating quantum development, AI shortens the window the adversary is wagering against, which makes the harvest a progressively better bet each year [1]. These two effects compound, since better triage raises the value of what is stored while a shorter wait lowers the cost of storing it. Security commentary through 2025 and 2026 has repeatedly identified HNDL as the largest board-level quantum risk for these reasons [12].

2.4. The Agentic AI Complication

A parallel development compounds the exposure. The enterprise shift from conversational AI to agentic AI, meaning systems that not only generate text but act, decide, and execute against connected tools and data sources, expands the attack surface that quantum-era defenses must protect. When an organization grants autonomous agents' discretion over tools, databases, and external interfaces, every agent request becomes a potential vector, and the traditional human-centric perimeter model erodes further. Agentic systems also generate and consume long-lived secrets at machine speed, multiplying the volume of sensitive material in transit and at rest. The intersection with quantum risk is direct rather than incidental, because the long-lived secrets, credentials, model weights, and proprietary datasets that agentic systems routinely handle are among the assets most exposed to HNDL collection [12]. Securing the AI layer and migrating to quantum-resistant cryptography are therefore not separate initiatives to be sequenced independently. They are two aspects of a single architectural problem that must be reasoned about together, since a gap in either undermines the other. An organization that hardens its agent permissions but leaves its communications on quantum-vulnerable cryptography has closed one door and left another open.

2.5. Interpreting the Timeline Debate

The structure of the timeline debate itself merits attention, because experts often misread disagreement as grounds for complacency. The range of credible estimates for a cryptographically relevant quantum computer remains wide, spanning

in various assessments from the late 2020s to the 2040s, and that dispersion is real [1]. The error lies in inferring from the breadth of the range that the prudent response is to wait for it to narrow. Three considerations argue against that inference. First, the relevant decision variable for most organizations is not the expected arrival date but the lower tail of the distribution, so that a credible, non-trivial probability of a CRQC by the early 2030s places at risk any data that must remain confidential through that period. Second, the HNDL dynamic means that protective action taken after a CRQC arrives is too late for data already collected, which lowers the option value of waiting well below its apparent level [12]. Third, the direction of revision matters more than the absolute estimate, and successive analyses have moved the expected timeline earlier, as the reduction in qubit requirements and the doubling of expert probability estimates since 2022 both illustrate [1] [2]. A planner should therefore weight the trajectory rather than the current midpoint. The appropriate reading of the debate is not that experts disagree, and time remains, but that experts disagree, tail risk is material, and the cost of late action is asymmetric.

This asymmetry is what the AI accelerant intensifies. In a world where quantum progress depended only on the unaided pace of hardware engineering, the wide estimate range might justify a measured wait-and-see posture. In a world where a general-purpose accelerant is applied to the field's defining bottlenecks, the same range should be read as a distribution whose mass is migrating toward the near term [1]. The distinction matters for governance, because a board that treats the timeline as static will underinvest relative to one that treats it as shifting. The convergence does not resolve the uncertainty about the arrival of a CRQC, often termed Q-Day. Instead, it shifts the responsible planning posture under that uncertainty, moving the reasonable default from deferral toward preparation. Uncertainty, in this framing, is an argument for action rather than against it.

3. The Defensive Vector: Artificial Intelligence as the Engine of Quantum Resilience

If the convergence were purely offensive, the appropriate posture would be defensive triage, but it is not purely offensive. The same coupling of AI and quantum computing operates in the defensive direction, and for most organizations, defensive applications are nearer-term and more immediately actionable than the offensive threat. Two applications deserve particular attention. The first is the maturing QML research program for threat detection, which is promising but not yet operational. The second is AI's immediate operational role in executing cryptographic migration, which is available now. This section treats them in turn, distinguishing carefully between demonstrated capability and prospective capability so that the strategic recommendations in Section 5 rest on the former where possible.

3.1. Quantum Machine Learning for Threat Detection

QML is a research-stage methodology that combines quantum computing and

machine learning to explore whether quantum resources, including superposition, entanglement, and quantum parallelism, can offer computational advantages over classical methods for specific learning tasks. In cybersecurity, those tasks principally include intrusion and anomaly detection, malware and botnet classification, and encrypted-traffic analytics [3]. Whether QML delivers a meaningful advantage in any given application depends on the specific task, dataset characteristics, the quality of the classical baseline used for comparison, and the hardware or simulation environment in which the quantum model is evaluated, conditions that vary significantly across the published literature and that limit direct comparison of reported results. The results reviewed in this section should therefore be read as evidence of exploratory potential rather than of settled, generalizable advantage.

In cybersecurity, the principal application areas are intrusion and anomaly detection, malware and botnet classification, and encrypted-traffic analytics [3]. A 2025 taxonomy of the field maps the relevant techniques, including quantum neural networks (QNNs), quantum support vector machines (QSVMs), variational quantum circuits, and quantum generative adversarial networks, across supervised, unsupervised, and generative learning paradigms and onto these core security tasks [3]. The taxonomy is useful because it organizes a literature that has grown quickly and unevenly, allowing practitioners to locate a given method within the broader landscape. It also clarifies which security tasks have attracted sustained research attention, and which remain underexplored. That mapping provides the structure for reading the empirical results discussed below.

The empirical results warrant precise attribution, as the figures originate in distinct primary studies with specific experimental conditions. In intrusion detection, Abreu *et al.* (2025) evaluated the QuantumNetSec architecture, a hybrid quantum-classical intrusion-detection system, on standard network benchmark datasets including NSL-KDD and UNSW-NB15, reporting superior performance in both binary and multiclass classification compared to classical-only baselines under simulated quantum conditions [4]. The taxonomy by Sai *et al.* (2025) synthesizes a range of quantum neural network approaches that have achieved reported accuracies of 96 to 97 percent in intrusion-detection experiments, and summarizes one quantum-assisted system reaching 96.8 percent accuracy against 92.3 percent for comparable classical approaches; however, these figures aggregate across heterogeneous experimental settings and should be read as indicative of the range rather than as standardized benchmarks [3] [5]. In zero-day detection, Shukla *et al.* (2025) report an F1 score of 0.94 for a quantum support vector machine approach against 0.87 for classical baselines, evaluated on a curated malicious-intrusion dataset using quantum circuit simulation rather than physical quantum hardware [6]. Detection latency improvements, including the reduction from 320 milliseconds to 110 milliseconds cited in the taxonomy, have been reported in network anomaly-detection experiments on Internet of Things (IoT) traffic, such as Deepak (2025), using quantum-classical hybrid methods and sim-

ulation-based evaluation [5]. None of the primary studies cited here evaluated their approaches on live production networks or physical quantum processors at scale; all are simulation- or emulation-based, and generalization to operational environments remains to be established.

These results must be read with appropriate caution, and the literature that reports them is candid about the limits. Quantum models currently lack the interpretability and transparency often required in cybersecurity decision-making, which constrains their use in settings where analysts must justify an alert [5]. Many studies rely on limited datasets, which raises legitimate questions about generalization across diverse network environments and attack types [3] [5]. Hardware noise, interoperability, and cost remain substantial barriers to operational adoption, and none of these is likely to resolve quickly [3]. A fair characterization is that QML for cybersecurity is a maturing research program rather than a deployable capability. Its strategic significance lies in the parallel it draws to the offensive timeline, since the defensive and offensive races run on the same hardware curve. As quantum capability matures, it will be available not only to adversaries decrypting harvested data but also to defenders modeling, classifying, and anticipating threats, and organizations familiar with QML methods will be positioned to absorb that capability as it arrives.

3.2. AI-Driven Migration: Solving the Execution Problem

AI's more immediate defensive contribution is less visible but more decisive. The central difficulty of the post-quantum transition is not the cryptography, because NIST has finalized the algorithms and, for most purposes, the question of what to migrate toward is settled [7]-[9]. The difficulty is execution: inventorying every cryptographic dependency across a heterogeneous enterprise, assessing exposure, sequencing migration in dependency order, managing governance and compliance, and tracking progress across thousands of systems and applications. This is an operational problem of large combinatorial scale, embedded in legacy environments that no human team could fully map by hand. Machine learning is well suited to this class of problem because pattern discovery across vast, irregular data is a core strength of the technology. Framing migration as an execution problem rather than a cryptographic one is therefore not a rhetorical move but a description of where the effort actually concentrates.

Enterprise cryptographic discovery tools available as of 2026 range in sophistication from rule-based scanners that identify specific algorithm signatures in codebases and network traffic, to platforms that augment this scanning with machine-learning techniques to infer cryptographic usage from behavioral patterns and indirect indicators, for example, classifying unknown encrypted flows by traffic fingerprinting, or flagging library call patterns consistent with deprecated algorithm use without direct code inspection. Vendors including IBM (through its Quantum Safe capabilities), Entrust (its PKI discovery tooling), and several NIST National Cybersecurity Excellence Partnership participants have published docu-

mentation describing automated cryptographic inventory functions for enterprise environments. Publicly available literature offers limited independent evaluation of these platforms' accuracy, coverage, and false-positive rates in heterogeneous production environments, and practitioners should treat vendor-reported discovery rates with the same scrutiny applied to any security tool claim. The literature does not always clearly distinguish between rule-based discovery, which reliably identifies known algorithm identifiers, and ML-augmented discovery, which attempts to infer cryptographic exposure from indirect signals and whose precision is highly environment-dependent. Both categories are in active commercial development as of 2026, and their comparative maturity varies by use case.

The strategic implication is favorable and easily overlooked. The organizations most exposed to the quantum timeline are large, complex enterprises with the most extensive cryptographic footprints, and those are the environments where AI-driven discovery and migration tooling delivers the greatest leverage. The relationship between exposure and the remedy's effectiveness is therefore aligned rather than opposed, which is an unusual and welcome feature of the problem. This defensive application does not depend on speculative future quantum hardware. It is deployable today against a migration burden that is already mandatory under the standards and deadlines examined in Section 4.

Table 1. AI as both the quantum threat accelerator and the quantum defense enabler.

Dimension	Offensive impact	Defensive impact	Organizational implication
Quantum computing	AI accelerates quantum error correction, qubit optimization, and algorithm development, reducing time toward a CRQC.	AI supports quantum hardware optimization and security analytics.	Organizations must assume quantum timelines continue to compress.
Cryptography	AI increases the urgency of Harvest Now, Decrypt Later (HNDL) attacks.	AI identifies vulnerable cryptographic assets requiring migration.	Immediate cryptographic inventory becomes essential.
Threat detection	Adversaries can automate targeting, prioritization, and attack optimization.	Quantum Machine Learning (QML) improves anomaly detection, malware detection, and intrusion detection.	Security operations become increasingly AI-driven.
Enterprise security	AI expands attack surfaces through agentic AI and autonomous systems.	AI enables enterprise-scale migration planning and cryptographic discovery.	Security architecture must integrate AI and post-quantum planning simultaneously.
Strategic outcome	Accelerated offensive capability.	Accelerated defensive capability.	AI becomes the catalyst for both quantum risk and quantum resilience.

Table 1 summarizes the article's central argument that artificial intelligence functions as a dual-force multiplier within the quantum cybersecurity landscape. On the offensive side, AI accelerates advances in quantum computing by improving error correction, algorithm optimization, and hardware efficiency, thereby reducing the estimated timeline toward a CRQC. At the same time, AI strengthens cybersecurity defenses through QML, automated cryptographic discovery, enterprise-wide asset identification, and post-quantum migration planning. The table

demonstrates that AI is not solely a source of cyber risk but also the primary technological enabler for achieving enterprise-scale quantum resilience, emphasizing that organizations must address both dimensions concurrently rather than as independent initiatives.

3.3. From Inventory to Sequenced Migration

Inventory is necessary but not sufficient, and the harder operational question is the order in which an organization migrates the dependencies it has discovered. Cryptographic systems are deeply interdependent, since certificates chain to roots, protocols negotiate algorithms dynamically, and applications inherit cryptographic behavior from shared libraries. Migration performed in the wrong order can introduce interoperability failures as damaging as the vulnerability being remediated. Three sequencing principles help structure the work. The first is exposure-weighted prioritization, under which systems handling data with the longest confidentiality requirements or the highest harvest value migrate first, consistent with the data-lifetime trigger described in Section 5. The second is dependency-aware ordering, under which foundational components such as certificate authorities, key-management infrastructure, and widely shared cryptographic libraries migrate before the systems that depend on them, so as to avoid repeated rework. The third is reversibility, under which each migration step is stageable and, where feasible, reversible, so that an interoperability failure can be contained rather than propagated. Together these principles convert a large undifferentiated backlog into an ordered program with a defensible rationale for each step.

This is where AI-driven tooling delivers value beyond initial discovery. Mapping the dependency graph across thousands of systems, identifying the foundational nodes whose migration unblocks the most downstream work, and re-evaluating the sequence as the environment changes are analytic tasks well matched to machine learning. A static inventory captures the environment at a moment, but real environments change continuously as systems are added, retired, and re-configured. An AI-maintained inventory can track that drift and flag newly introduced dependencies before they accumulate. The combination of an AI-maintained cryptographic inventory and an AI-informed sequencing model converts what would otherwise be an intractable manual program into a managed, prioritized, and auditable migration. In practice, this is the difference between a multi-year drift and a deadline-driven plan.

4. The 2026 Regulatory Inflection

The convergence is unfolding against a regulatory backdrop that, in a short span, has removed the principal justifications for inaction. In August 2024, NIST finalized its first three post-quantum cryptographic standards, namely FIPS 203 for key encapsulation, FIPS 204 for digital signatures, and FIPS 205 for a stateless hash-based signature scheme, all intended to replace the quantum-vulnerable RSA and ECC [7]-[9]. These are Federal Information Processing Standards, which

are mandatory for federal systems and widely adopted by organizations beyond the federal government. In March 2025, NIST selected HQC as a backup key-encapsulation mechanism, adding redundancy against the possibility that vulnerabilities might later emerge in the primary lattice-based schemes [10]. That choice is instructive, because it signals that the transition is not a single terminal replacement but an ongoing capability requiring continued vigilance. Standardizing a backup before the primary schemes are widely deployed reflects a deliberate hedging strategy. It also foreshadows the argument for cryptographic agility developed in Section 5.

The compliance timelines are now concrete rather than aspirational. NIST's guidance in Interagency Report 8547 describes phasing out quantum-vulnerable algorithms after 2030 and disallowing them entirely by 2035, with high-risk systems expected to transition earlier [10]. The National Security Agency's Commercial National Security Algorithm Suite 2.0 (CNSA 2.0) requires new national security systems to support quantum-safe algorithms beginning January 1, 2027, and cascades quantum-safe requirements through the defense industrial base through procurement [16]. For organizations operating in the defense and critical-infrastructure sectors, these are not distant targets but near-term procurement and compliance realities that govern current acquisition decisions. A supplier bidding on national security work in 2026 must already account for the 2027 milestone in its product roadmap. The deadlines therefore exert influence well before their nominal dates, shaping present decisions.

The broader institutional environment reflects a corresponding shift from awareness to coordination. An industry coalition declared 2026 the Year of Quantum Security, marked by January 2026 events featuring senior officials from the Federal Bureau of Investigation, NIST, and CISA [1] [17]. Market signals point in the same direction, since the post-quantum cryptography market, valued at roughly 1.9 billion United States dollars in 2025, is projected to reach 12.4 billion by 2035 [18]. The Boston Consulting Group has cautioned that organizations beginning migration in 2030 will already be too late, given the time required to inventory and transition complex systems [19]. These signals are mutually reinforcing, since regulatory deadlines drive market demand, and market growth in turn accelerates the tooling that makes compliance feasible. When the AI-driven compression of the quantum timeline is layered onto these defined regulatory horizons, the prudent planning assumption is not that the window is comfortably open but that it is narrower than headline estimates imply.

A Sector Lens: Critical Infrastructure and Healthcare

The convergence does not bear uniformly across sectors, and critical-infrastructure operators, healthcare among them, occupy a position of particular exposure. Three characteristics compound their risk. The first is data longevity: medical records, patient identifiers, and clinical data carry confidentiality requirements measured in decades, placing them squarely within the HNDL threat envelope [12].

For such data, the relevant question is not whether a CRQC arrives within five years but whether it arrives within the lifetime of a patient record, a horizon that spans most credible timeline estimates. The second characteristic is infrastructure heterogeneity, since healthcare environments combine modern systems with long-lived embedded devices and legacy clinical technology that cannot be patched on conventional cycles. This produces the sprawling, opaque cryptographic footprint that manual inventory cannot address and that AI-driven discovery is designed to map. The third characteristic is regulatory cascade, since providers and their technology suppliers connected to federal systems inherit post-quantum obligations through the same procurement mechanisms that bind the defense industrial base [16]. The CNSA 2.0 and NIST timelines therefore reach further into the private sector than is sometimes appreciated.

For these operators, the strategic framework in Section 5 applies with heightened force. The data-lifetime trigger becomes a near-certainty rather than a conservative assumption once record-retention requirements are considered, because clinical records must remain confidential far longer than any plausible deferral of quantum risk would allow. Given infrastructure heterogeneity, AI-driven inventory shifts from a convenience to a practical necessity, since no manual process can keep pace with the diversity of devices in a large clinical environment. Cryptographic agility likewise becomes a precondition for managing devices that will remain in service across multiple standards cycles, rather than an architectural preference. The same logic extends to any sector, including finance, energy, and government, in which data confidentiality requirements are long, infrastructure is heterogeneous, and regulatory obligations cascade through supply chains. This convergence is a general phenomenon, but it is most urgent where long-lived data meets complex, hard-to-inventory infrastructure. Healthcare simply presents that combination in an especially concentrated form.

5. A Five-Part Strategic Framework for Security Leaders

Translating the convergence into action does not require an organization to resolve every uncertainty about when a cryptographically relevant quantum computer will arrive. It requires accepting that the timeline is uncertain in a direction that favors urgency and acting on the elements already within reach. Five priorities follow from the foregoing analysis, ordered from the conceptual reframing that should govern prioritization to the concrete planning discipline that should govern execution. The five are complementary rather than sequential, since an organization can and should pursue them in parallel. Each is stated as a principle and then given an operational interpretation, so that the framework can be applied rather than merely endorsed. Taken together, they describe a posture of disciplined urgency calibrated to the evidence presented in Sections 2 through 4.

Figure 2 presents a practical five-stage framework for organizational quantum readiness that translates strategic recommendations into an actionable implementation model. The circular sequence emphasizes that quantum readiness is a con-

tinuous governance process rather than a one-time technology migration. Beginning with identifying data requiring long-term confidentiality, organizations move through an AI-driven cryptographic inventory, secure AI infrastructure, implement cryptographic agility, and align migration activities with regulatory deadlines set by NIST and CNSA 2.0. The cyclical design reinforces that enterprise quantum preparedness requires continuous assessment, modernization, and adaptation as quantum technologies, standards, and threat landscapes continue to evolve.

ENTERPRISE QUANTUM READINESS FRAMEWORK

A Five-Step Strategic Roadmap for Post-Quantum Security



Figure 2. Enterprise quantum readiness framework.

5.1. Trigger on Data Lifetime, Not on Q-Day

The operationally relevant question is not when quantum computers will break encryption but how long a given class of data must remain confidential. Any data whose confidentiality requirement extends past the early-2030s deprecation horizon is already exposed to HNDL collection and should be prioritized for quantum-resistant protection immediately [10] [12]. Reframing prioritization around data lifetime converts an unanswerable forecasting question into a tractable inventory-and-classification exercise that an organization can begin today. This reframing also has an organizational benefit, since data-retention requirements are typically already documented for compliance purposes and can be repurposed to drive migration priority. The practical first step is to cross-reference existing re-

tention schedules against cryptographic exposure. Data that must be retained for decades and is currently protected by quantum-vulnerable cryptography defines the initial migration cohort.

5.2. Inventory before Migrating, and Use AI to Do It

Migration cannot proceed rationally without a complete and current map of where vulnerable cryptography resides. At enterprise scale, manual inventory is infeasible, and AI-enhanced discovery tooling is the practical path to a comprehensive, continuously updated picture of cryptographic dependencies. Treat inventory as a standing capability rather than a one-time project, because both the environment and the standards will continue to evolve. A point-in-time inventory begins to decay the moment it is completed, as new systems are deployed and old ones reconfigured. Continuous, automated discovery addresses that decay by design. The output of this capability feeds directly into the sequencing discipline described in Section 3.3, making inventory the foundation on which the rest of the program rests.

5.3. Secure the AI Layer as Part of the Quantum Transition

Model weights, training data, and the secrets exposed by agentic systems are high-value, long-lived assets that fall within the harvest-now crosshairs [12]. Plan post-quantum protection for AI infrastructure alongside, not after, the migration of traditional data-at-rest and data-in-transit. The access discretion granted to autonomous agents should be reasoned about as part of the same threat model, since an over-privileged agent operating over quantum-vulnerable channels compounds two risks at once. Treating AI security and quantum migration as one program avoids the common failure of hardening an AI deployment against present threats while leaving its communications exposed to future decryption. This integration is easier to achieve at design time than to retrofit later. Organizations building agentic systems now can embed quantum-resistant protections before those systems reach production scale.

5.4. Adopt Cryptographic Agility Rather than Point Fixes

Selecting HQC as a backup, because the primary schemes could still reveal weaknesses, underscores that the transition is not a single rip-and-replace event but a durable architectural property [10]. Systems should be designed so that cryptographic algorithms can be substituted with minimal disruption as standards evolve, vulnerabilities surface, and new backups are standardized. Cryptographic agility is the difference between a one-time migration cost and a recurring crisis, because an agile system absorbs the next algorithm change as routine maintenance rather than as an emergency. Achieving agility requires abstracting cryptographic operations behind well-defined interfaces rather than hard-coding algorithms into applications. This is an established software-engineering discipline applied to a new purpose. Organizations that build agility into the current migration will find sub-

sequent transitions far less costly.

5.5. Reverse-Plan from Binding Deadlines

For organizations touching national security systems or the defense industrial base, the CNSA 2.0 milestone of 2027 and the broader NIST horizon of 2030 to 2035 are compliance facts, not planning assumptions [10] [16]. Migration roadmaps should be reverse-engineered from these dates, sequencing high-risk and long-lived systems first so the organization arrives at each deadline with its most exposed assets already protected. Reverse planning also surfaces resource constraints early, since working backward from a fixed date reveals whether the available staff and budget can complete the required work in time. Where they cannot, the gap can be addressed by adjusting scope, adding resources, or escalating to leadership well before the deadline becomes acute. This discipline turns abstract deadlines into concrete quarterly milestones. It is the planning counterpart to the technical sequencing described earlier and closes the gap between strategy and execution.

Table 2. Five strategic priorities for Quantum-AI cybersecurity readiness.

Strategic priority	Objective	Expected organizational benefit
Trigger on data lifetime	Protect long-lived sensitive information before CRQC arrival.	Reduces harvest now, decrypt later exposure.
AI-driven cryptographic inventory	Discover cryptographic dependencies automatically.	Improves visibility across enterprise environments.
Secure the AI layer	Protect AI models, data, credentials, and agentic systems.	Prevents AI from becoming a quantum vulnerability.
Cryptographic agility	Design systems capable of replacing algorithms rapidly.	Supports future standards without major redesign.
Reverse plan from regulatory deadlines	Align migration with NIST and CNSA 2.0 milestones.	Improves compliance, governance, and organizational readiness.

Table 2 translates the conceptual discussion into five practical strategic priorities for organizations preparing for the post-quantum era. Rather than waiting for the arrival of a CRQC, the framework emphasizes proactive risk management through data lifetime prioritization, AI-assisted cryptographic inventory, protection of AI infrastructure, cryptographic agility, and compliance-driven migration planning. Collectively, these priorities provide security leaders with an operational roadmap for reducing long-term cryptographic exposure while improving organizational resilience against rapidly evolving quantum-enabled threats

6. Global Investment Landscape and Geopolitical Competition

The AI-quantum convergence described in the preceding sections is unfolding inside a geopolitical competition that shapes both the pace of the threat and the adequacy of the response. Three state actors, the United States, China, and the

European Union, have each designated quantum computing a defined national strategic priority, but the character of their commitments differs in ways that bear directly on the timeline and defensive posture arguments developed above. This competitive context is not supplementary to the analysis; it shapes the timeline's uncertainty in terms of urgency, not merely uncertainty.

U.S. federal quantum investment, coordinated through the National Quantum Initiative and distributed across the Department of Energy, NIST, the National Science Foundation, and the intelligence community, exceeds four billion dollars annually as of 2025 [20]. Private sector investment from IBM, Google, Microsoft, and Amazon has substantially amplified that commitment, producing hardware milestones, including Google's Willow processor, that have driven the downward revision in qubit-requirement estimates discussed in Section 2.2. China's posture differs in a structurally significant way. Its Fourteenth Five-Year Plan designates quantum information as a frontier technology requiring national breakthrough, and its military-civil fusion framework ensures that advances in civilian quantum research are, by policy design, available to the People's Liberation Army without the institutional barriers that separate commercial and defense research in Western systems [21]. This integration is directly relevant to the HNDL threat: an adversary whose quantum research program is structurally connected to its signal's intelligence apparatus is one whose harvest-now operations and decryption timeline are governed by a unified strategic calculus. The European Quantum Flagship and the EuroQCI quantum communication infrastructure initiative reflect a third model, emphasizing defensive investment in quantum-secured communications networks and multinational coordination rather than competition on the offensive quantum computing timeline [22].

The investment landscape also introduces two complications that the strategic framework in Section 5 does not fully address. The first is supply chain concentration. Quantum hardware, specifically the dilution refrigerators, cryogenic control electronics, and specialized materials that superconducting qubit systems require, is manufactured by a small number of firms in a small number of countries, with rare earth inputs subject to the same Chinese supply chain leverage that the broader technology sector has identified as a national security concern [23]. Organizations planning QKD infrastructure or quantum-enhanced security capabilities should treat hardware supply chain risk as a planning variable alongside cryptographic algorithm selection. The second complication is standards fragmentation. China has developed indigenous post-quantum cryptographic standards, including the ShangMi algorithm suite, rather than adopting NIST's algorithms, and is actively participating in ISO and ITU standards processes to shape the global standards environment [21]. The result will not be a single unified post-quantum cryptographic order; at minimum, it will be two parallel ecosystems, each with different algorithms, key parameters, and governance timelines. Multinational organizations and supply chain partners with significant cross-border operations should treat this divergence as a present-day architecture decision, not a future

complication. **Table 3** summarizes the comparative investment postures of the three principal actors.

Table 3. National quantum investment postures: comparative overview.

Dimension	United states	China	European union
Governance model	Federal coordination (NQI) + private sector	State-directed; military-civil fusion	Multinational public coordination (Flagship)
Primary strategic objective	Maintain technology leadership; PQC standards	Quantum superiority. sovereign capability	Secure communications infrastructure
Military-civilian integration	Separated by institutional design	Unified by policy	Separated, NATO coordination
Cryptographic standards posture	NIST FIPS 203/204/205; CNSA 2.0	Indigenous (ShangMi). parallel ITU/ISO participation	ETSI QKD; ENISA guidance; NIST-aligned
Key defensive initiative	PQC migration mandates; CISA guidance	QKD backbone (Beijing-Shanghai)	EuroQCI quantum communication network

Note. NQI = national quantum initiative; MCF = military-civil fusion; PQC = post-quantum cryptography; CNSA 2.0 = commercial national security algorithm suite 2.0; ENISA = EU agency for cybersecurity.

For security leaders, three planning implications follow. First, treat the HNDL threat described in Section 2.3 as the expected behavior of a well-resourced, institutionally motivated adversary rather than a speculative worst case—the investment architecture described here incentivizes and enables systematic harvest-now collection. Second, post-quantum migration vendor selection should incorporate hardware supply chain scrutiny alongside algorithm evaluation, applying the same rigor to quantum-relevant components that mature programs apply to semiconductor and network equipment sourcing. Third, organizations with cross-border operations in Chinese-governed environments should plan for dual-stack cryptographic implementations rather than assuming a single global migration path. The five strategic priorities in Section 5 remain the appropriate operational response, but they are more urgently actionable and more accurately scoped when understood within the competitive context this section describes.

7. Discussion

The analysis presented across the preceding sections converges on a single claim: the security implications of quantum computing and artificial intelligence can no longer be assessed independently, because the two technologies now shape one another's trajectories. This section clarifies what that claim means for theory, practice, and the assumptions that have governed enterprise planning to date, and it situates the argument against the objections most likely to be raised.

The first implication concerns how organizations reason about time. Conventional planning has treated the arrival of a cryptographically relevant quantum computer as an exogenous event to be awaited, with preparation calibrated to a forecast horizon. The evidence assembled here suggests that this framing misstates the problem. When the resource requirement for breaking RSA-2048 falls from

roughly twenty million qubits toward one million, and when expert probability estimates roughly double over a two-year span, the timeline reveals itself as a moving quantity rather than a fixed date [1] [2]. The direction of that movement matters more than its current midpoint, because a variable being pushed consistently earlier by a general-purpose accelerant behaves differently from one drifting on the unaided pace of hardware engineering. The practical consequence is that planning anchored to a single expected arrival date is fragile in the very way that harms the organization.

The second implication concerns the locus of risk. The HNDL dynamic relocates quantum exposure from the future into the present for any data with a long confidentiality horizon [11] [12]. This relocation is easy to underweight because it produces no visible incident at the moment of compromise; the breach is silent, and its disclosure is deferred until decryption becomes feasible. Conventional incident-response frameworks, which assume that compromise becomes known within a bounded period, are poorly suited to a threat whose defining feature is latency. The organizations most affected are precisely those whose data must remain confidential for decades, and the analysis in Section 4.1 indicates that healthcare and other critical-infrastructure sectors sit at the sharp end of this exposure [12].

The third implication is more constructive and most likely to be overlooked. The same coupling that intensifies the threat also supplies the defense, and for most organizations, defensive applications are nearer at hand than the threat is imminent. QML remains a maturing research program constrained by limited datasets, hardware noise, and an interpretability gap, and its reported benchmarks should be read as evidence of potential rather than settled performance [3] [5]. The AI-driven discovery and sequencing of cryptographic dependencies, by contrast, is deployable now against a migration burden that is already mandatory [7]-[10]. A favorable and underappreciated feature of the problem is that exposure and remedy are aligned: the large, complex enterprises most exposed to the quantum timeline are the environments in which AI-driven tooling delivers the greatest leverage.

These implications should be weighed against several objections. The most common is that expert disagreement about the timeline justifies deferral. The analysis in Section 2.5 resists that inference on the grounds that the relevant decision variable is the lower tail of the distribution rather than its midpoint, that the option value of waiting is reduced by the retroactive nature of HNDL, and that the trajectory of successive revisions is itself informative [1] [12]. A second objection holds that the migration burden can be addressed later, once the threat is nearer. This objection underestimates the time required to inventory, sequence, and transition cryptography across a heterogeneous enterprise, a program that industry analysis suggests cannot be compressed into the final years before a deadline [18]. A third objection questions whether QML will deliver operational value at all. The argument does not depend on resolving that question, because AI's

immediate defensive contribution lies in migration rather than in QML, and the strategic case for QML rests on the structural point that offensive and defensive capability advance along the same hardware curve.

Taken together, these considerations reframe the appropriate posture under uncertainty. The responsible response to an uncertain but earlier-trending timeline, a present-tense data-exposure threat, and binding regulatory deadlines is neither alarmism nor deferral, but the disciplined urgency articulated in the strategic framework of Section 5. That framework is robust to the residual uncertainty about when a cryptographically relevant machine will arrive, because each of its five commitments is justified across the plausible range of timelines rather than contingent on any single forecast. This robustness distinguishes an actionable program from a speculative one, and it is the practical payoff of treating the convergence, rather than either technology in isolation, as the unit of analysis.

The geopolitical investment landscape examined in Section 7 sharpens two of the three implications developed above in ways that the technical and regulatory analysis alone does not fully capture. The first concerns how organizations should reason about the HNDL threat. The analysis in Section 2.3 established that harvest-now collection is occurring and that its economic logic is sound for any adversary who expects decryption capability within the useful lifetime of collected data. The investment picture adds a structural dimension to that argument: the actors most likely to conduct systematic HNDL collection are the same actors who have made state-directed quantum development a national priority, who have institutionally integrated their civilian quantum research programs with military and intelligence functions, and who have every strategic incentive to begin collection now against a decryption capability they are actively working to develop [12] [21]. This is not a speculative threat profile constructed from worst-case assumptions. It is the expected behavior of an adversary whose institutional architecture is designed for exactly this kind of long-horizon, capability-anticipatory collection. Security leaders who have treated HNDL as a theoretical concern warranting eventual attention should recalibrate that posture in light of the investment and integration picture the geopolitical analysis reveals [22].

The second implication concerns the scope of the migration problem. Section 5 describes post-quantum migration as a relatively clean transition: NIST has finalized the algorithms, the compliance deadlines are defined, and the strategic framework follows logically from those fixed points. The geopolitical context complicates that picture in two respects. China's development of indigenous post-quantum cryptographic standards, rather than adopting NIST's, means organizations with significant cross-border operations, supply chain relationships, or technology partnerships in Chinese-governed environments will not face a single global migration path. They will face at minimum two parallel cryptographic ecosystems with different algorithms, different governance authorities, and different timelines, requiring dual-stack implementations and ongoing interoperability management that the compliance-deadline framing of Section 4 does not antici-

pate. Separately, supply-chain concentration in quantum hardware components, dilution refrigerators, cryogenic electronics, and rare earth inputs introduces a physical dependency risk for organizations planning QKD infrastructure or quantum-enhanced security capabilities that sits outside the cryptographic algorithm migration discussion entirely. Mature post-quantum planning programs should address both complications explicitly rather than treating the NIST compliance horizon as the full scope of the problem. The earlier objection that the migration burden can be addressed later, once the threat is nearer, is least defensible for the organizations most affected by these complications, since dual-stack implementation and supply chain assessment are programs that cannot be compressed into the final months before a deadline [23].

8. Conclusion and Directions for Further Research

The habit of treating quantum computing and artificial intelligence as separate security horizons has outlived its analytical usefulness. AI is compressing the quantum threat timeline even as it furnishes a viable means of defending against it [1]-[3]. The HNDL dynamic has already converted a prospective future risk into a present exposure for long-lived data [12] [13]. The 2026 regulatory environment has finalized the standards and fixed the deadlines that govern the response [7]-[10] [16]. The defining cybersecurity challenge of the coming decade resides in the interaction of these technologies rather than in either one considered alone. Recognizing that interaction is the precondition for a proportionate response.

For security leaders, this analysis implies a posture that is neither alarmism nor the comfortable deferral distant estimates have long licensed. It is disciplined urgency, expressed in five concrete commitments: act on the data whose lifetime already exposes it, deploy AI against the migration burden now, secure the AI layer within the same effort, adopt cryptographic agility as a design principle, and reverse-plan from deadlines that are no longer hypothetical. These commitments are within the reach of any organization willing to begin. None depends on resolving uncertainty about Q-Day, since each is justified across the plausible range of timelines. That robustness to uncertainty is what makes the framework actionable rather than speculative. The organizations that adopt it will be prepared regardless of when the threat matures.

This analysis carries limitations that bound its claims. As a conceptual synthesis, it inherits the uncertainties of its sources, since estimates of the quantum timeline remain contested; the QML benchmarks reported in the literature are drawn from constrained experimental settings whose generalization to production environments is not established [2] [3] [5]. The article also reflects the state of the standards and threat environment as of mid-2026, and a domain moving this quickly will revise some specifics. Readers should therefore treat individual figures as indicative rather than definitive. These limitations qualify the particulars without disturbing the central argument, which depends not on any single figure but on the structural fact of the convergence and the asymmetric cost of late response.

The direction of the evidence is more robust than any individual data point.

This framing suggests several directions for further research. The empirical literature on QML for cybersecurity, though promising, remains constrained by limited datasets, hardware noise, and a persistent interpretability gap, and rigorous evaluation on diverse, operationally representative data is needed before QML can move from research program to deployable capability [3] [5]. Governance of agentic AI under post-quantum threat models is largely unexplored and warrants dedicated study, particularly how autonomous agents should manage long-lived secrets under the assumption of eventual decryption. The economics of HNDL, including how adversaries value, triage, and prioritize harvested data and how AI alters that calculus, remains under-studied relative to its strategic importance [12]. Each of these questions sits at the intersection this article has described, which suggests that the convergence is not only a practical challenge but a research agenda. What is already clear is that organizations need not wait for that agenda to be completed before they act.

The geopolitical dimension of the AI-quantum convergence reinforces the conclusion's central argument while expanding its scope. The disciplined urgency the framework calls for is not only a response to an uncertain technical timeline; it is a response to a competitive environment in which adversaries have made strategic commitments to quantum capability, are applying the same AI accelerants to their development programs that compress the timeline for everyone, and are institutionally structured to conduct harvest-now collection against data that Western organizations are still protecting with algorithms those adversaries are actively working to break. Standards fragmentation between U.S. and Chinese cryptographic ecosystems, and supply chain concentration in the physical components that quantum security infrastructure requires, mean that the migration problem is larger and less tractable than the compliance-deadline framing alone suggests [21] [22]. Organizations that scope their post-quantum programs to the NIST horizon and no further will find themselves underprepared for a threat environment that does not respect the boundaries of any single regulatory framework. The five commitments in Section 5 remain the right foundation, but they are most actionable, and most urgently actionable, when understood as moves in a competition already underway rather than preparations for a contingency still pending [23].

Conflicts of Interest

The author declares no conflict of interest regarding the publication of this paper.

References

- [1] SecurityWeek. (2026) Cyber Insights 2026: Quantum Computing and the Potential Synergy with Advanced AI. SecurityWeek.
<https://www.securityweek.com/cyber-insights-2026-quantum-computing-and-the-potential-synergy-with-advanced-ai/>

- [2] Riverlane (2026) Quantum Error Correction: 2025 Trends and 2026 Predictions. Riverlane.
<https://www.riverlane.com/blog/quantum-error-correction-our-2025-trends-and-2026-predictions>
- [3] Sai, S., Kanadia, M. and Chamola, V. (2025) Quantum Machine Learning for Cyber-Security: A Taxonomy and Future Directions. arXiv: 2512.15286.
<https://arxiv.org/abs/2512.15286>
- [4] Abreu, D., Moura, D., Esteve Rothenberg, C. and Abelém, A. (2025) QuantumNetSec: Quantum Machine Learning for Network Security. *International Journal of Network Management*, **35**, e70018. <https://doi.org/10.1002/nem.70018>
- [5] Kavitha, S.S., Dhanush Sagar, M.D. and Deepak, S. (2025) Internet of Things Network Intrusion Detection System Using Quantum and Classical Machine Learning. *IET Quantum Communication*, **6**, e70020. <https://doi.org/10.1049/qtc2.70020>
- [6] Shukla, A.K. (2025) Safeguarding Networks from Malicious Intrusions Using Quantum Machine Intelligence. *Software: Practice and Experience*, **55**, 1676-1695.
<https://doi.org/10.1002/spe.70004>
- [7] National Institute of Standards and Technology (2024) Module-Lattice-Based Key-Encapsulation Mechanism Standard (FIPS 203). U.S. Department of Commerce.
<https://doi.org/10.6028/NIST.FIPS.203>
- [8] National Institute of Standards and Technology (2024) Module-Lattice-Based Digital Signature Standard (FIPS 204). U.S. Department of Commerce.
<https://doi.org/10.6028/NIST.FIPS.204>
- [9] National Institute of Standards and Technology (2024) Stateless Hash-Based Digital Signature Standard (FIPS 205). U.S. Department of Commerce.
<https://doi.org/10.6028/NIST.FIPS.205>
- [10] National Institute of Standards and Technology (2025) Transition to Post-Quantum Cryptography Standards (NIST IR 8547). U.S. Department of Commerce.
<https://doi.org/10.6028/NIST.IR.8547.ipd>
- [11] Cloud Security Alliance AI Safety Initiative (2026) Harvest Now, Decrypt Later: Quantum Risk to AI Infrastructure. Cloud Security Alliance.
<https://labs.cloudsecurityalliance.org/research/ai-infrastructure-post-quantum-harvest-now-decrypt-later-v1/>
- [12] ISC2 (2026) Harvest Now, Decrypt Later. International Information System Security Certification Consortium.
<https://www.isc2.org/insights/2026/05/harvest-now-decrypt-later?queryID=7e5d7ff453374e366389678314feba0e>
- [13] Shor, P.W (1994) Algorithms for Quantum Computation: Discrete Logarithms and Factoring. *Proceedings 35th Annual Symposium on Foundations of Computer Science*, Santa Fe, 20-22 November 1994, 124-134.
<https://doi.org/10.1109/sfcs.1994.365700>
- [14] Ghosh, T. and Susanna, R.C. (2026) Toward Intelligent Fault-Tolerant Quantum Computing: A Review of Google's Willow Processor and AI Integration. In 2026 *IEEE International Conference for Convergence in Computing Technology (I3CTCON)* (pp. 1-6). IEEE. <https://doi.org/10.1109/I3CTCON68242.2026.11507170>
- [15] Cybersecurity and Infrastructure Security Agency (2022) Prepare for a New Cryptographic Standard to Protect against Future Quantum-Based Threats. U.S. Department of Homeland Security.
<https://www.cisa.gov/news-events/alerts/2022/07/05/prepare-new-cryptographic->

- [standard-protect-against-future-quantum-based-threats](#)
- [16] National Security Agency (2022) Commercial National Security Algorithm Suite 2.0 (CNSA 2.0). U.S. Department of Defense.
https://media.defense.gov/2025/May/30/2003728741/-1/-1/0/CSA_CNSA_2.0_ALGORITHMS.PDF
 - [17] The Quantum Insider (2026) Why 2026 Matters for Quantum Security. The Quantum Insider.
<https://thequantuminsider.com/2026/04/28/why-2026-matters-quantum-security/>
 - [18] Aaronson, S., Childs, A.M., Farhi, E., Harrow, A.W. and Sanders, B.C. (2026) Future of Quantum Computing. *Quantum Machine Intelligence*, **8**, 3.
<https://doi.org/10.1007/s42484-026-00351-2>
 - [19] Biswas, S. (2026) AI-Driven Quantum-Resistant Security Framework for Accelerating the Migration of US Cloud and IOT Systems to NIST Post-Quantum Cryptography Standards. *American Journal of Advanced Technology and Engineering Solutions*, **16**, 1-44. <https://doi.org/10.63125/8pwsq720>
 - [20] Yulchiev, K. (2025) The Growth of Quantum Computing: Opportunities and Strategic Advantages for the United States. *IMRAS*, **8**, 148-156.
<https://doi.org/10.5281/zenodo.14988841>.
 - [21] Kania, E.B. and Costello, J.K. (2018) Quantum Hegemony? China's Ambitions and the Challenge to U.S. Innovation Leadership. Center for a New American Security.
https://s3.us-east-1.amazonaws.com/files.cnas.org/documents/CNASReport-Quantum-Tech_FINAL.pdf
 - [22] European Commission. (2023). European Quantum Communication Infrastructure (EuroQCI) Initiative.
<https://digital-strategy.ec.europa.eu/en/policies/european-quantum-communication-infrastructure-euroqci>
 - [23] Congressional Research Service (2024) Quantum Computing: Concepts, Current State, and Considerations for Congress (CRS Report R47685).
https://www.congress.gov/crs_external_products/R/PDF/R47685/R47685.2.pdf