CSAI Foundation | Cloud Security Alliance

Harvest Now, Decrypt Later: Quantum Risk to AI Infrastructure

Post-Quantum Readiness as an Enterprise-Wide Systemic Risk

2026-05-18

AI-assisted Rapid Research

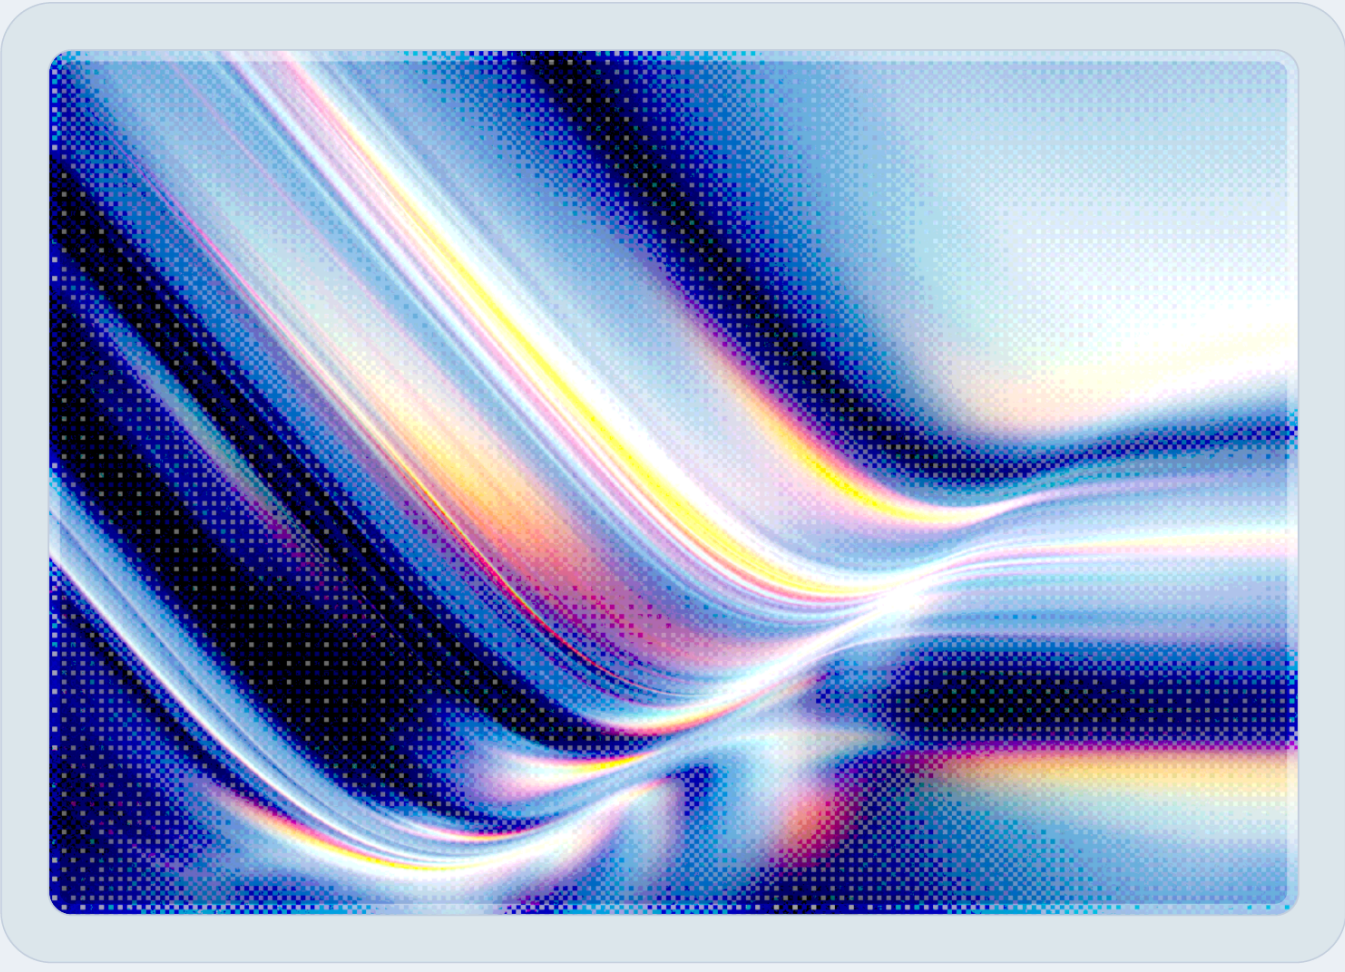

CSAI

CSA cloud security alliance®

Table of Contents

Executive Summary

The quantum computing threat to enterprise cryptography is not a forecast — it is an ongoing operation. The attack pattern known as "Harvest Now, Decrypt Later" (HNDL) describes a strategy, well-documented by Western intelligence agencies and national cybersecurity authorities, in which adversaries systematically intercept and archive encrypted data today, holding it in reserve against the day when cryptographically relevant quantum computers (CRQCs) can be used to break that encryption retroactively. For organizations building and operating AI infrastructure, this creates a category of risk that is qualitatively different from conventional cybersecurity threats: the breach may already have occurred, the data may already be in adversarial hands, and the organization may not know it until years from now.

AI infrastructure carries unusually high HNDL exposure. Model weights represent years of compute investment and proprietary research; training datasets can contain sensitive personal, financial, or medical data subject to long-retention compliance requirements; inference endpoints serve as real-time intelligence feeds into enterprise decision-making; and the multi-agent communication architectures now proliferating across enterprise workflows create dense networks of internal trust boundaries, each typically protected by TLS — a protocol whose underlying public-key cryptography is fundamentally vulnerable to a sufficiently powerful quantum computer.

The regulatory response has advanced materially since August 2024, when NIST published FIPS 203, 204, and 205 — the first finalized post-quantum cryptographic standards — and NIST Interagency Report 8547 (Initial Public Draft) established a formal deprecation timeline calling for the disallowance of RSA and elliptic curve cryptography across NIST standards by 2035 [1][2]. The NSA's Commercial National Security Algorithm Suite 2.0 (CNSA 2.0) imposes binding deadlines on national security systems and cascades quantum-safe requirements through the defense industrial base, with new acquisitions for National Security Systems expected to support CNSA 2.0 algorithms beginning January 1, 2027 [3]. Enterprise timelines for full PQC migration range from five to fifteen or more years depending on organization size and cryptographic complexity [4]. For large enterprises beginning migration in 2026, a meaningful risk window exists in which quantum-capable decryption may be feasible before migration is complete.

This paper examines the HNDL threat as it specifically bears on AI infrastructure, maps the regulatory and standards landscape now in force, and presents a phased migration framework aligned to CSA's AI Controls Matrix and Cloud Controls Matrix.

Introduction: A Breach That Happens in Reverse

Cybersecurity has long operated on the assumption that breach detection and incident response can be organized around the moment of compromise. HNDL undermines that assumption in a fundamental way. When an adversary intercepts TLS-encrypted traffic between an AI training cluster and a cloud storage backend, the adversary has collected data that is, in the present moment, unreadable. Nothing alerts. No integrity check fails. No anomaly is flagged. The encrypted payload is archived and held, potentially indefinitely, awaiting a future key capable of unlocking it.

The concept is not new — intelligence agencies have engaged in long-range signals collection for decades — but the combination of two converging developments has elevated HNDL from a theoretical concern to an active and urgent operational risk. First, quantum computing hardware has advanced substantially. IBM has outlined a roadmap to deliver a fault-tolerant quantum system using approximately 10,000 physical qubits by 2029 [28], and Quantinuum, in collaboration with Microsoft, has demonstrated logical qubit error rates 800 times lower than corresponding physical qubit benchmarks [29]. Forrester Research's 2026 assessment explicitly characterized Q-Day — the arrival of a cryptographically relevant quantum computer — as a plausible risk by 2030 [6]. Second, three papers published between May 2025 and March 2026 reduced the estimated quantum resources required to break RSA-2048 from approximately twenty million qubits to fewer than one million, with some architectural proposals suggesting the attack might be achievable with as few as 100,000 qubits [7]. These are not proofs of imminent capability; they are demonstrations of a trajectory that has consistently surprised in the direction of accelerating risk.

Against this backdrop, the U.S. Department of Homeland Security, the UK National Cyber Security Centre, the European Union Agency for Cybersecurity, and the Australian Cyber Security Centre have all issued post-quantum guidance premised on the assumption that sophisticated adversaries are currently collecting and storing sensitive encrypted traffic for future decryption [30][31][32][33]. A 2025 Federal Reserve working paper specifically examined HNDL risks to distributed ledger networks and discussed how encrypted data in transit today could be exposed to retroactive decryption as quantum capabilities mature [9]. The National Endowment for Democracy has documented China's development of frontier technologies — including quantum computing and quantum communication infrastructure — as instruments of a data-centric strategic technology program with global security implications [10].

For enterprise security professionals, the practical implication is direct: any data transmitted over classical public-key cryptography today that will retain sensitivity in 2030 or beyond is potentially at risk. This is not a future consideration. The harvest is occurring now.

The HNDL Threat Landscape

Who Is Collecting and What They Are Targeting

Open-source and governmental assessments describe HNDL activity primarily in the context of nation-state intelligence operations. The architecture of the attack is straightforward: actors with access to internet transit infrastructure, undersea cables, or network choke points intercept bulk encrypted traffic and archive it in large-scale data stores [8]. The marginal cost of storing encrypted ciphertext is low; the potential future value — if quantum decryption becomes feasible — is high.

High-value HNDL targets generally share two characteristics. They involve data that retains strategic, commercial, or intelligence value over a multi-year horizon, and they rely on classical public-key cryptography for protection in transit or at rest. Government communications, intellectual property, financial records, health data, and strategic business documents clearly meet the first criterion. Modern AI infrastructure meets both.

The targeting rationale for AI systems is substantial. A trained foundation model represents a concentrated artifact of enormous compute investment — in many cases billions of dollars in GPU cycles, curated data, and engineering labor. The model's weights, if exfiltrated, could be used to replicate capabilities that took years to develop, to reverse-engineer training data for sensitive personal information, or to conduct adversarial research that undermines the model's intended use. Training datasets themselves may contain health records, financial transactions, or communications that were originally encrypted under the assumption that they would remain confidential indefinitely. Inference-time communications between enterprise AI systems and downstream consumers can reveal strategic priorities, operational patterns, and competitive intelligence.

The 2025 NSA Cybersecurity Information Sheet on AI data security, co-signed by CISA, the FBI, and allied agencies, specifically calls for organizations running AI systems to adopt quantum-resistant cryptographic standards, citing the combination of high-value AI assets and the long operational lives of deployed systems as factors that elevate HNDL exposure [11].

The Quantum Timeline

Estimates of when a cryptographically relevant quantum computer will arrive have consistently narrowed. National Security Memorandum-10, signed in 2022, required federal agencies to complete post-quantum migration by 2035 — implicitly acknowledging that a CRQC could plausibly exist before that date. NSA, CISA, and NIST issued a joint advisory urging organizations to begin post-quantum migration immediately,

given the possibility of Q-Day within the current decade [34]. Forrester's 2026 assessment, among the most recent authoritative commercial research on the question, described practical quantum computing as feasible within five years and treated Q-Day as a live risk scenario for enterprise planning [6].

The significance of these timeline estimates lies not in their precision — no one can predict a technological breakthrough with certainty — but in their implication for data with a long sensitivity horizon. An organization that transmits sensitive AI model parameters or proprietary training data under RSA-2048 today, and where a CRQC arrives in 2030, has roughly four years from now to complete migration before that data could in principle be decrypted. For large enterprises, where realistic PQC migration timelines range from twelve to fifteen or more years, the math is unfavorable [4].

AI Infrastructure as a High-Value HNDL Target

The Unique Exposure Profile of AI Systems

AI infrastructure differs from conventional enterprise IT in several characteristics that amplify HNDL risk. Understanding these differences is necessary to prioritize migration efforts effectively.

First, AI artifacts have unusually long sensitivity horizons. A trained model deployed in a regulated industry may be in production for three to seven years. The training data used to produce it may be legally required to be retained for a decade or more. Model weights, if they incorporate proprietary techniques or trade secrets, may represent competitive advantages that remain relevant far beyond typical software life cycles. All of this means that data encrypted today must remain confidential through a period in which quantum decryption may become feasible.

Second, AI pipelines involve multiple cryptographic trust boundaries that each represent an interception opportunity. Data ingestion from source systems, preprocessing and transformation in cloud storage, transfer to training clusters, gradient synchronization in distributed training, model checkpointing and registry storage, distribution to inference endpoints, and API communication with downstream applications — each of these transitions involves data in motion, typically protected by TLS or similar protocols whose key exchange mechanisms rely on RSA or elliptic curve Diffie-Hellman [12]. Each hop is a potential collection point for HNDL operations.

Third, the proliferation of agentic AI architectures has created dense networks of internal machine-to-machine communication that operate largely outside traditional network security monitoring. When an AI agent invokes a tool, queries an external data source, or communicates with a peer agent using protocols such as Model Context Protocol (MCP) or agent-to-agent APIs, that communication is protected by TLS. In a multi-agent deployment, dozens or hundreds of such channels may exist, each a potential point of

interception [13]. The encrypted inter-agent traffic generated by enterprise AI workflows today — describing tool calls, data queries, decision rationales, and intermediate results — could reveal substantial operational intelligence if decrypted in the future.

Fourth, AI model supply chains depend on cryptographic signatures for integrity and provenance. The emerging ecosystem of model signing and attestation — including the OpenSSF Model Signing specification, NVIDIA's integration with NGC, and Sigstore-based provenance chains — relies on digital signature schemes that will be compromised by a CRQC [14][15]. An adversary who harvests signed model artifacts today could, in the future, forge provenance claims by breaking the signatures on legitimate artifacts, potentially enabling supply chain attacks that insert backdoored weights into production systems while presenting valid historical attestations.

Mapping the AI Cryptographic Attack Surface

The following table summarizes the key cryptographic exposure points in a representative enterprise AI infrastructure and their relative HNDL sensitivity.

Component	Typical Cryptographic Protection	HNDL Sensitivity	Primary Risk
Training data in cloud storage	AES-256 at rest; RSA/ECC key wrapping	High	IP and PII exposure when keys are broken
Data pipelines in transit	TLS 1.3 (ECDHE key exchange)	High	Long-lived sensitive data intercepted in transit
Distributed training gradients	TLS between worker nodes	Medium-High	IP leakage; reconstruction of training data
Model checkpoints / weights	TLS in transit; AES-256 at rest	Very High	Core IP; years of compute investment
Model registry / artifact store	TLS + code signing (ECDSA/RSA)	Very High	Provenance forgery; backdoor insertion
Inference API endpoints	TLS 1.3	High	Reveals usage patterns, queries, outputs

Component	Typical Cryptographic Protection	HNDL Sensitivity	Primary Risk
Agent-to-agent communication	TLS (per connection)	Medium-High	Operational intelligence; decision rationale
Authentication (API keys, tokens)	RSA/ECDSA; JWTs	High	Future credential compromise
Federated learning aggregation	TLS + homomorphic/differential privacy	Medium	Gradient reconstruction at scale

The central observation from this mapping is that symmetric cryptography (AES-256) protecting data at rest is not directly broken by Shor's algorithm — AES-256 is considered quantum-resistant with appropriate key lengths [1]. The exposure is concentrated in the public-key mechanisms used for key exchange, digital signatures, and authentication. This has an important practical implication: the confidentiality protection of data at rest under AES-256 is not itself the primary vulnerability. The vulnerability lies in the key management and transport layer — specifically, the RSA and ECC mechanisms by which symmetric keys are wrapped, exchanged, and protected. If those mechanisms are broken, data encrypted under them can be decrypted regardless of the strength of the underlying symmetric cipher.

The Regulatory and Standards Landscape

NIST's Post-Quantum Framework

The foundational regulatory development in post-quantum cryptography is NIST's publication of FIPS 203, 204, and 205 in August 2024 — the first finalized post-quantum cryptographic standards [1][16]. FIPS 203 specifies ML-KEM (Module-Lattice-Based Key-Encapsulation Mechanism, derived from the CRYSTALS-Kyber algorithm), which provides quantum-resistant key encapsulation for key exchange scenarios. FIPS 204 specifies ML-DSA (Module-Lattice-Based Digital Signature Algorithm, derived from CRYSTALS-Dilithium), the primary standard for quantum-resistant digital signatures. FIPS 205 specifies SLH-DSA (Stateless Hash-Based Digital Signature Standard), a hash-based backup signature scheme using different mathematical assumptions. Together, these three standards provide the foundational building blocks for quantum-resistant cryptographic infrastructure across enterprise systems.

NIST Interagency Report 8547, published as an Initial Public Draft in November 2024, establishes the proposed deprecation timeline [2]. RSA-2048 and ECC P-256 — the most commonly deployed public-key algorithms in enterprise infrastructure — are designated for deprecation by 2030, meaning they should not be used in new deployments after that date. By 2035, all quantum-vulnerable public-key algorithms will be formally disallowed under NIST standards, regardless of key length. The practical implication is that any AI infrastructure deployed today that is expected to operate through 2030 or beyond should incorporate PQC migration planning now.

NSA CNSA 2.0 and the 2027 Inflection Point

The NSA's Commercial National Security Algorithm Suite 2.0 creates binding obligations for national security systems and cascades through the defense industrial base. Starting January 1, 2027, NSA expects new acquisitions for National Security Systems to support CNSA 2.0 algorithms — requiring ML-KEM-1024, ML-DSA-87, AES-256, and SHA-384/512 at minimum [3]. Equipment that cannot support CNSA 2.0 must be phased out by December 31, 2030, with full mandatory CNSA 2.0 adoption across all National Security Systems required by December 31, 2031. Compliance with these requirements cascades downstream to defense contractors, sub-tier suppliers, and commercial vendors whose products interact with national security systems. Organizations developing AI systems for federal or defense applications must incorporate this timeline into their architecture and procurement decisions immediately.

CISA and Federal Civilian Guidance

CISA has published a strategy for migrating to automated post-quantum cryptography discovery and inventory (ACDI) tools, establishing a framework for how federal civilian agencies and critical infrastructure operators should approach the discovery and prioritization phases of PQC migration [17]. The strategy emphasizes that cryptographic visibility — a complete, continuously updated inventory of where and how cryptography is used across enterprise systems — is a prerequisite for effective migration. For AI organizations, this has particular relevance because AI pipelines typically involve dozens of interconnected systems, each with its own cryptographic footprint, spanning cloud provider services, on-premises infrastructure, third-party APIs, and managed services.

International Regulatory Alignment

The UK NCSC's migration guidance defines three phases culminating in complete PQC adoption by 2035, with Phase 1 — cryptographic discovery and migration planning — targeted for completion by 2028 [31]. The European Union Agency for Cybersecurity has published technical guidelines aligned to NIST's standards [32]. Australia's ASD/ACSC has issued guidance recommending that organizations cease reliance on traditional asymmetric cryptography for new deployments by the end of 2030 and complete detailed

transition plans by the end of 2026 [33]. For multinational AI organizations, this regulatory convergence simplifies compliance in one sense — the algorithms required (ML-KEM, ML-DSA, SLH-DSA) are consistent across jurisdictions — but requires attention to jurisdiction-specific timelines and reporting obligations.

The Migration Challenge

Cryptographic Debt and Discovery

The single most commonly cited barrier to enterprise PQC migration is not technical — it is informational. Organizations that have operated IT infrastructure for a decade or more have accumulated substantial cryptographic debt: encryption, signing, and authentication mechanisms embedded in legacy applications, firmware, hardware security modules, cloud service configurations, SaaS APIs, and custom code, with little or no centralized visibility into what exists where. Before migration can begin, organizations must know what they have.

NIST's NCCoE Cryptographic Discovery program has validated agent-based discovery tools capable of cataloging cryptographic usage across network traffic, file systems, software packages, and database systems [17]. Commercial tools from vendors including Keyfactor (which acquired InfoSec Global in 2025), AppViewX, and others provide automated cryptographic inventory capabilities aligned to the NCCoE's methodology. The concept of a Cryptographic Bill of Materials (CBOM) — analogous to a software bill of materials — provides a structured artifact for documenting and tracking cryptographic dependencies throughout the supply chain [18]. For AI organizations managing complex pipelines with many vendors, establishing CBOM practices for AI artifacts and their dependencies is a foundational first step.

Performance and Compatibility Considerations

Post-quantum algorithms carry different performance and size characteristics compared to RSA and ECC. ML-KEM encapsulations and decapsulations are computationally efficient and comparable in performance to existing key exchange mechanisms, but ML-KEM public keys are approximately 1,184 bytes (versus 32 bytes for a P-256 public key), and ciphertexts are approximately 1,088 bytes — implications that matter for high-throughput API gateways and constrained edge inference devices [1]. ML-DSA signatures are larger still, with signatures reaching approximately 3,293 bytes for ML-DSA-65 [35].

These size increases have direct operational implications for AI infrastructure components that rely on certificate-based authentication, TLS handshakes under load, or signed attestations transmitted alongside model artifacts. Inference endpoints serving thousands of requests per second may require TLS library upgrades, hardware offloading, or architectural changes to handle increased handshake overhead. IoT-

adjacent AI deployment scenarios — such as on-device inference in manufacturing or healthcare — may involve hardware that cannot accommodate larger post-quantum key and signature formats without replacement.

A hybrid cryptographic approach, layering PQC algorithms alongside classical algorithms in a combined construction, is recommended during the transition period [19]. Hybrid approaches provide forward secrecy against quantum attacks on harvested traffic while maintaining backward compatibility with existing systems. IETF hybrid key exchange specifications for TLS 1.3 provide a standardized mechanism for this approach. Meta's published PQC migration framework, released in April 2026, is one of the most detailed public accounts of how a large technology organization is managing this hybrid transition at scale, and is a useful reference for enterprises planning similar efforts [20].

The Migration Gap

Research published in 2025 and 2026 consistently identifies a structural risk arising from the intersection of accelerating quantum timelines and extended enterprise migration timelines. A study published in Computers (MDPI) in 2025 estimated realistic migration timelines of five to seven years for small enterprises, eight to twelve years for medium enterprises, and twelve to fifteen or more years for large enterprises [4]. If Q-Day arrives by 2030 — a scenario increasingly cited by Forrester, NIST, and NSA — organizations beginning PQC migration in 2026 may face a window of several years during which significant portions of their infrastructure remain quantum-vulnerable. Traffic encrypted during that window would be susceptible to retroactive decryption.

For AI organizations, this migration gap analysis has a concrete implication: the highest-sensitivity AI data flows should be prioritized for early PQC adoption, regardless of where they fall in a comprehensive migration roadmap. This means that protecting model weight distribution channels, training data pipelines, and model registry signing should be treated as early-phase priorities, not items deferred to later in a multi-year migration program.

The talent constraint compounds the timeline challenge. There is a documented global shortage of cryptography engineers with expertise in post-quantum algorithms, and enterprises are competing with government agencies and defense contractors for a limited pool of specialists [4]. Training existing security engineering teams and establishing crypto-agility as an architectural standard — the ability to swap cryptographic algorithms without rewriting application logic — is preferable to relying on external expertise for a multi-year migration.

A Framework for Post-Quantum Readiness in AI Environments

Post-quantum readiness for AI infrastructure should be organized around five sequential phases. These are not strictly linear — discovery and risk assessment should continue throughout the migration lifecycle — but they provide a practical structure for program management.

Phase 1: Discover and Inventory

The foundation of any PQC migration program is cryptographic visibility. Organizations should deploy automated cryptographic discovery tools to enumerate every instance of RSA, ECC, and other quantum-vulnerable public-key algorithms across their AI infrastructure. This inventory should span network traffic (intercepting TLS handshakes to identify negotiated cipher suites), software packages (scanning code and dependencies for cryptographic library usage), certificate stores, hardware security modules, cloud KMS configurations, and SaaS API integrations. The output should be a structured CBOM for each AI system, with sufficient detail to identify algorithm type, key length, certificate expiry, and the system component or software layer responsible for each cryptographic dependency.

For AI-specific inventory, particular attention should be paid to model signing and attestation mechanisms, training pipeline encryption configurations, inter-service TLS configurations in agentic architectures, and the cryptographic dependencies of ML frameworks and data orchestration tools (TensorFlow, PyTorch, Kubeflow, Airflow, and their associated dependencies all use cryptographic libraries for transport and authentication that require audit).

Phase 2: Classify and Prioritize

Not all cryptographic usage carries the same HNDL risk. Organizations should classify discovered cryptographic instances by two dimensions: the sensitivity of the data protected, and the expected sensitivity horizon — the length of time that data will need to remain confidential. Data that retains its sensitivity for more than five years — a conservative estimate of the potential Q-Day window — should be treated as high HNDL exposure. For AI organizations, this typically includes model weights, proprietary training datasets containing personal or financial data, long-term API keys and service credentials, and strategic communications about AI capability development.

Risk prioritization should also account for the criticality of the system performing the cryptographic operation. An inference endpoint serving external customers is both a higher-value HNDL target and a higher-risk migration target due to availability requirements. Internal development systems may be lower-

value but also lower-risk to migrate. The prioritization matrix should inform migration sequencing, not just risk reporting.

Phase 3: Implement Hybrid Cryptography

Migration to post-quantum algorithms should begin with hybrid implementations that layer ML-KEM over existing ECDHE key exchange in TLS connections, and ML-DSA alongside ECDSA in signing workflows. Hybrid constructions are specified in IETF drafts for TLS 1.3 and provide immediate HNDL protection — any traffic protected by a hybrid TLS connection is quantum-resistant even if the classical component is later broken — while maintaining interoperability with systems that have not yet been upgraded [21].

For AI model registries and supply chain signing, migration to ML-DSA-based signatures for model artifacts, SBOMs, and attestations should be an early priority given the forward-looking implications: a model signed today with RSA or ECDSA may need to be trusted in 2031, by which point the signature could in principle be forged. Adopting ML-DSA for new model artifacts now, and re-signing critical existing artifacts, reduces the risk of future provenance attacks. The OpenSSF Model Signing specification supports post-quantum algorithm selection and provides a framework for this migration [14].

For data at rest, the primary remediation focus should be on key management: replacing RSA-based key wrapping with ML-KEM-based key encapsulation in cloud KMS configurations and hardware security modules. Existing AES-256-encrypted data does not need to be re-encrypted — it is the key protection mechanism that requires updating.

Phase 4: Validate and Certify

Post-quantum migration introduces new failure modes that differ from classical cryptographic implementation errors. Organizations should validate PQC implementations through formal testing against NIST test vectors, ensure correct parameter selection (ML-KEM-768 or ML-KEM-1024, not lower parameter sets, for high-sensitivity applications), and confirm that hybrid constructions are implemented correctly so that security is bounded by the stronger of the two algorithms. For organizations building AI systems intended for use with national security customers, CNSA 2.0 validation requirements include National Information Assurance Partnership (NIAP) or CSfC program validation — a significant lead-time consideration that should be built into product roadmaps now.

Performance testing under post-quantum load is particularly important for high-throughput AI inference serving. The increased TLS handshake size and signature verification overhead of PQC algorithms should be characterized under production-representative load, and any necessary infrastructure changes — load balancer upgrades, TLS offload hardware, connection multiplexing — should be addressed before cutover.

Phase 5: Monitor and Maintain Crypto-Agility

Post-quantum migration is not a one-time event. NIST's post-quantum standardization process is ongoing; a fourth standard (FN-DSA, based on FALCON) has been approved, and additional algorithms are under evaluation [1]. The threat landscape may evolve in ways that affect algorithm confidence. Crypto-agility — the architectural property of being able to replace cryptographic algorithms without rewriting business logic — is the long-term design principle that reduces the cost of future transitions.

For AI organizations, crypto-agility means abstracting cryptographic operations behind configurable interfaces in ML framework integrations, training infrastructure, and inference serving code; maintaining updated CBOMs for all AI artifacts and dependencies; and monitoring the NIST Post-Quantum Cryptography project, IETF standardization activity, and vendor library releases for algorithm updates that should be incorporated into production systems.

Conclusions and Recommendations

HNDL is not a hypothetical future scenario; it is an active collection strategy being executed by nation-state adversaries against enterprise networks today. For AI organizations, the combination of high-value proprietary assets, complex multi-hop cryptographic infrastructure, and long data sensitivity horizons creates an exposure profile that warrants urgent attention. The regulatory framework — NIST FIPS 203/204/205, NIST IR 8547, NSA CNSA 2.0, and aligned international guidance — has reached a level of maturity that provides a clear technical direction. The question for enterprises is not what to do, but how to execute a migration at organizational scale within a narrowing window.

The following recommendations are organized by time horizon.

Immediate Actions (0–12 months)

Organizations should initiate cryptographic discovery across all AI systems to establish a baseline CBOM, focusing first on data pipelines, model registries, and inter-service communication in agentic architectures. Model signing workflows for new AI artifacts should be migrated to ML-DSA at the earliest feasible point to ensure forward-looking provenance integrity. TLS configurations for the highest-sensitivity AI data channels — model weight distribution, training data ingestion from external sources, external inference APIs — should be upgraded to hybrid TLS supporting ML-KEM key exchange. CNSA 2.0 compliance timelines should be incorporated into product roadmaps for any AI system intended for federal, defense, or critical infrastructure customers.

Short-Term Actions (12–36 months)

Organizations should complete cryptographic inventory across all AI infrastructure, establish CBOM practices as part of AI supply chain management, and begin hybrid TLS rollout across internal service-to-service communication in AI platforms. Key management infrastructure — cloud KMS, HSMs, PKI — should be upgraded to support ML-KEM key encapsulation for symmetric key protection. Existing model artifacts with long expected deployment lives should be re-signed with ML-DSA. Security teams should receive training in post-quantum algorithm selection, implementation, and validation.

Strategic Considerations (36 months and beyond)

The long-term objective is complete migration away from quantum-vulnerable public-key algorithms in all AI infrastructure, aligned to the NIST IR 8547 deprecation timeline. This requires treating crypto-agility as a first-class architectural requirement for all new AI systems, incorporating PQC requirements into vendor procurement standards and supply chain due diligence, and maintaining continuous cryptographic visibility as AI infrastructure evolves. Organizations should monitor NIST post-quantum standardization activity for updates to the algorithm landscape, and should be prepared to incorporate additional or revised standards as the field matures.

CSA Resource Alignment

This whitepaper connects directly to several CSA frameworks and publications that together provide the governance and technical architecture for post-quantum AI security.

AICM (AI Controls Matrix)

The AICM's cryptography and data protection control domains address key management, encryption in transit and at rest, and supply chain integrity — each of which requires post-quantum migration planning. AI organizations should map their PQC roadmaps to AICM control domains to assess readiness and identify gaps. The AICM's shared responsibility model is particularly relevant for AI organizations using cloud-managed KMS services: the customer retains responsibility for the key management policies and algorithm selection even when the underlying hardware is managed by a cloud provider.

CCM (Cloud Controls Matrix)

CSA's Quantum-Safe Security Governance with the Cloud Controls Matrix provides specific CCM control mappings for quantum threats, including controls on encryption, key management, and change management that should be updated to reflect PQC requirements [22]. The CCM's encryption controls should be interpreted against NIST FIPS 203/204/205 as the current standard of practice.

CSA Post-Quantum Cryptography Publications

The CSA Quantum-Safe Security Working Group has produced a substantial body of work directly applicable to enterprise PQC migration, including *A Practitioner's Guide to Post-Quantum Cryptography* [23], *The State of Post-Quantum Cryptography* [24], *Practical Preparations for the Post-Quantum World* [25], and *Applied Quantum Safe Security* [26]. These publications provide the technical foundation for migration programs and should be consulted alongside this whitepaper. The CSA Labs Q-Day Clock research provides enterprise migration timeline analysis aligned to current quantum computing projections [27].

MAESTRO (Multi-Agent and Agentic AI Threat Modeling)

MAESTRO's threat modeling framework for agentic AI should be extended to include HNDL as an explicit threat class in multi-agent architectures. The inter-agent communication channels cataloged in MAESTRO threat models are exactly the TLS-protected data flows that carry HNDL exposure, and PQC migration of those channels should be integrated into agentic architecture security reviews.

STAR (Security Trust Assurance and Risk)

CSA STAR assessments for AI products and services should incorporate post-quantum readiness as an evaluation dimension, including cryptographic inventory completion, PQC adoption status, and CNSA 2.0 compliance posture for applicable customers.

References

[1] NIST. "Federal Information Processing Standard 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard." NIST, August 2024.

[2] NIST. "IR 8547 (Initial Public Draft): Transition to Post-Quantum Cryptography Standards." NIST, November 2024.

[3] QuSecure. "CNSA 2.0 Explained: PQC Requirements, Timelines, and Federal Impact." QuSecure, 2025.

[4] MDPI Computers. "Enterprise Migration to Post-Quantum Cryptography: Timeline Analysis and Strategic Frameworks." MDPI, 2025.

[5] Palo Alto Networks. "What Is Q-Day, and How Far Away Is It—Really?." Palo Alto Networks, 2025.

[6] Forrester Research. "Practical Quantum Computing By 2030 Is Likely — And So Is Q-Day." Forrester, 2026.

[7] The Quantum Insider. "What Is 'Harvest Now, Decrypt Later' and Why Should You Care?." The Quantum Insider, May 2026.

[8] Palo Alto Networks. "Harvest Now, Decrypt Later (HNDL): The Quantum-Era Threat." Palo Alto Networks, 2025.

[9] Federal Reserve. "'Harvest Now Decrypt Later': Examining Post-Quantum Cryptography and the Data Privacy Risks for Distributed Ledger Networks." Federal Reserve Finance and Economics Discussion Series, 2025.

[10] National Endowment for Democracy. "Data-Centric Authoritarianism: How China's Development of Frontier Technologies Could Globalize Repression." National Endowment for Democracy, 2025.

[11] NSA, CISA, FBI et al. "Joint Cybersecurity Information: AI Data Security." U.S. Department of Defense, May 2025.

[12] Security Boulevard. "The Future of AI Cyber Security: Why Quantum-Resistant Encryption Is Non-Negotiable." Security Boulevard, May 2026.

[13] Gopher Security. "Post-Quantum AI Infrastructure Security: Protecting MCP Deployments in 2026." Gopher Security, 2026.

[14] OpenSSF. "An Introduction to the OpenSSF Model Signing (OMS) Specification." Open Source Security Foundation, June 2025.

[15] NVIDIA. "Bringing Verifiable Trust to AI Models: Model Signing in NGC." NVIDIA Technical Blog, 2025.

[16] NIST. "NIST Releases First 3 Finalized Post-Quantum Encryption Standards." NIST News, August 2024.

[17] CISA. "Strategy for Migrating to Automated Post-Quantum Cryptography Discovery and Inventory Tools." CISA, 2024.

[18] ReversingLabs. "Accelerate PQC Migration: How to Leverage CBOMs for Cryptographic Asset Discovery." ReversingLabs, 2025.

[19] Fortinet. "Achieve Crypto-Agility for Quantum Readiness." Fortinet, 2025.

[20] Meta Engineering. "Post-Quantum Cryptography Migration at Meta: Framework, Lessons, and Takeaways." Engineering at Meta, April 2026.

[21] NIST NCCoE. "Migration to Post-Quantum Cryptography: Crypto-Agility Considerations." NCCoE, 2025.

[22] CSA. "Quantum-Safe Security Governance with the Cloud Controls Matrix." Cloud Security Alliance, 2024.

[23] CSA. "A Practitioner's Guide to Post-Quantum Cryptography." Cloud Security Alliance, 2025.

[24] CSA. "The State of Post-Quantum Cryptography." Cloud Security Alliance, 2018.

[25] CSA. "Practical Preparations for the Post-Quantum World." Cloud Security Alliance, 2021.

[26] CSA. "Applied Quantum Safe Security." Cloud Security Alliance, 2017.

[27] CSA Labs. "Q-Day Clock: Enterprise Post-Quantum Migration Imperative." CSA Lab Space, 2025.

[28] IBM. "IBM Lays Out Clear Path to Fault-Tolerant Quantum Computing." IBM Quantum Blog, June 2025.

[29] Microsoft. "Advancing Science: Microsoft and Quantinuum Demonstrate the Most Reliable Logical Qubits on Record with an Error Rate 800x Better than Physical Qubits." Microsoft Blog, April 2024.

[30] DHS. "Post-Quantum Cryptography." U.S. Department of Homeland Security, 2024.

[31] UK NCSC. "Timelines for Migration to Post-Quantum Cryptography." National Cyber Security Centre, 2025.

[32] ENISA. "Post-Quantum Cryptography: Current State and Quantum Mitigation." European Union Agency for Cybersecurity, 2022.

[33] ASD/ACSC. "Planning for Post-Quantum Cryptography." Australian Signals Directorate, 2025.

[34] NSA/CISA/NIST. "Post-Quantum Cryptography: CISA, NIST, and NSA Recommend How to Prepare Now." National Security Agency, August 2023.

[35] NIST. "Federal Information Processing Standard 204: Module-Lattice-Based Digital Signature Algorithm Standard." NIST, August 2024.