PLANCKCHRON

TRUSTWORTHY AI · POST-QUANTUM READY

Post-Quantum Readiness for Enterprise AI

A practical migration framework for the harvest-now, decrypt-later era.

White paper · 2026

Planckchron Inc. · Los Angeles, CA · www.planckchron.com

EXECUTIVE SUMMARY

The cryptography that protects today's AI systems has an expiry date

A cryptographically relevant quantum computer does not yet exist publicly — but the data your organization encrypts today can be captured today and decrypted later, once one does. For the long-lived secrets that sit at the heart of modern AI — training corpora, model weights, customer data, API keys, and signing keys — that future risk is a present-day exposure. This paper explains why AI estates are unusually vulnerable to the “harvest-now, decrypt-later” threat, summarizes the now-finalized post-quantum cryptography (PQC) standards, and lays out a five-step framework any enterprise can use to become post-quantum ready without halting delivery.

1. The threat: harvest now, decrypt later

Most public-key cryptography in use today — RSA and elliptic-curve schemes that secure TLS, VPNs, code signing, and stored secrets — rests on mathematical problems that a sufficiently powerful quantum computer could solve efficiently using Shor's algorithm. The timeline for such a machine is uncertain and debated. The risk, however, is not contingent on that timeline.

An adversary can intercept and store encrypted traffic and exfiltrated archives now, and simply wait. The moment quantum decryption becomes feasible, everything harvested in the interim is retroactively exposed. The relevant question for any security team is therefore not **“when will quantum computers arrive?”** but **“how long does my data need to stay confidential?”** Any secret whose required confidentiality lifetime extends into the 2030s is already, in effect, at risk.

2. Why AI systems are disproportionately exposed

AI estates concentrate exactly the kind of high-value, long-lived data that the harvest-now threat targets — and they do so across a sprawling, fast-moving supply chain:

- **Model weights and training data** are expensive, durable assets. A model trained today may remain commercially or operationally sensitive for a decade — well within the exposure window.
- **Data pipelines** move sensitive records between storage, feature stores, training clusters, and inference endpoints, multiplying the number of channels an attacker can harvest.
- **Long-lived secrets** — API keys, service credentials, and signing keys for model artifacts — are frequently rotated slowly, if at all.
- **Vendor and model sprawl** means most organizations cannot even enumerate every AI system, third-party API, and dependency that touches regulated data.

Step 3 — Architect for crypto-agility 2 — Tier: rank by data lifetime × harvest exposure; build a cryptographic inventory

The durable goal is not a single swap to one algorithm but the ability to change algorithms readily as standards and threats evolve. Abstract cryptographic choices behind well-defined interfaces, centralize key management, and remove hard-coded algorithm assumptions so future transitions are configuration changes rather than re-engineering efforts.

Step 4 — Migrate in phases, using hybrid modes

Adopt the standardized PQC algorithms highest-risk-tier first. Where supported, deploy hybrid modes that combine a classical and a post-quantum algorithm, so security holds even if either is later weakened, while interoperability is preserved. Sequence the rollout against the risk tiers from Step 2 and validate at each stage.

Step 5 — Assure continuously: governance and evidence

Treat readiness as an ongoing posture, not a one-time project. Continuously monitor for newly introduced cryptography and AI systems, track migration progress against the inventory, and maintain audit-ready evidence for regulators, customers, and counsel. Crypto-agility plus continuous assurance is what keeps an organization ready as the standards and the threat landscape keep moving.

5. How Planckchron approaches this

Planckchron is an independent advanced-technology company focused on the trust, security, and assurance layer for enterprise and government AI. Our methodology maps directly to the framework above:

- **Assure — visibility & governance:** establish the cryptographic and AI-system inventory and the risk picture (Steps 1–2).
- **Guard — runtime security:** protect AI systems in production with monitoring and audit logging that support continuous assurance (Step 5).
- **Post-Quantum Readiness:** map cryptographic exposure to the harvest-now threat and chart a phased, standards-aligned migration path (Steps 2–4).

Our work is model-agnostic and standards-aligned: we do not train frontier models, and we do not replace your cryptographic providers. We help you see your exposure clearly and migrate it in a defensible order.

Getting started

A post-quantum readiness assessment is the practical first move: a structured discovery and risk-tiering exercise that turns an abstract future threat into a prioritized, costed plan. To request one, visit **www.planckchron.com** or contact **ai@planckchron.com**.

About this paper. This white paper is a positioning and methodology document intended to help security and engineering leaders reason about post-quantum risk in AI systems. It describes publicly available standards and a general framework; it is not security, legal, or investment advice, and it does not represent a guarantee of any specific outcome. Standards and

government timelines referenced (NIST FIPS 203/204/205; NSA CNSA 2.0) are summarized as of 2026 and should be verified against the issuing bodies for the latest detail. Planckchron's products and services referenced here describe our methodology and current focus.