Harvest Now, Decrypt Later as a Balance-Sheet Risk: Pricing the Quantum Threat to Long-Dated Financial Instruments

ACM Global Tech

Abstract

The cryptography that secures a long-dated financial instrument is not an IT detail; it is an assumption embedded in the instrument's value, and that assumption has a *duration*. A "harvest now, decrypt later" (HNDL) adversary records ciphertext, public keys, and signatures today and breaks them once a cryptographically relevant quantum computer (CRQC) exists. For a 30-year bond, a perpetual stablecoin reserve, or a multi-decade custody record, the relevant question is not whether a CRQC arrives but whether it arrives *before the instrument matures*—and whether the issuer has migrated by then. This paper reframes the quantum threat for a Chief Risk Officer and CFO audience: as a *duration mismatch* between instrument tenor and cryptographic half-life, quantified through Michele Mosca's inequality $x + y > z$, and priced through a simple expected-loss model, $\mathbb{E}[\text{loss}] \approx \text{exposure} \times \mathbb{P}(\text{break before maturity})$. We show how migration shortens the effective *cryptographic duration* of an exposure, give a governance sequencing across custody, settlement, identity, and archives, and describe ACM Global Tech's crypto-agile answer—NIST-standardized post-quantum algorithms as native precompiles from genesis, with hybrid migration—built on Lux (blockchain, post-quantum cryptography, FHE), Hanzo (AI/data), and Zen (LLMs). All numbers in the risk model are explicitly *illustrative*; the framework, not the figures, is the deliverable.

1 The Risk, Stated Plainly

A cryptographic assumption protects a financial instrument only for as long as that assumption holds. The classical public-key schemes underpinning today's settlement rails, custody systems, and digital ledgers—RSA, ECDSA over secp256k1, Ed25519—rest on the hardness of integer factoring and discrete logarithms. Shor's algorithm solves both in polynomial time on a sufficiently large quantum computer, so a CRQC can recover a private key from its public key and forge signatures at will [4, 6].

The threat does not wait for that machine to exist. A patient adversary can *harvest now and decrypt later*: record ciphertext, public keys, and signatures off the wire and from public ledgers today, store them cheaply, and break them the day a CRQC is available [4, 8]. Three properties make this acute for finance:

- **The data is already exposed.** On a public blockchain, every public key and signature is permanently readable. There is no "recall"; what is recorded today is harvested by default.
- **The instruments are long-dated.** A 30-year government or corporate bond, a perpetual stablecoin reserve, a multi-decade pension or insurance liability, and a custody record that must survive an audit horizon all carry confidentiality and integrity requirements measured in *decades*.

- **The break is retroactive.** A forged signature against a key harvested in 2026 can be produced in, say, 2034 and presented as authentic. The integrity of a multi-decade record is only as strong as the weakest cryptography it ever relied on.

For a long-dated instrument, then, the cryptographic assumption is itself a *liability with a duration*. It does not appear on the balance sheet, but it behaves like one: a contingent obligation that comes due if the assumption fails before the instrument does.

2 Why This Is a Balance-Sheet Item, Not an IT Line Item

Risk managers already reason about *duration*: the sensitivity of an instrument's value to the horizon over which its assumptions must hold. The quantum threat is best understood the same way—as a mismatch between two durations.

Two durations. Let an instrument have *tenor* T_{mat}: the time until maturity or, for perpetuals and archives, the horizon over which integrity and confidentiality must hold. Let the protecting cryptography have a *cryptographic half-life*: the time until the underlying hard problem is expected to be broken at scale. When tenor exceeds cryptographic half-life, the instrument outlives its own security. This is a duration mismatch in the precise risk sense: the asset's assumptions are funded shorter than the asset itself, and the gap is a position the institution is short whether or not it has chosen to be.

Mosca's inequality. Michele Mosca formalized the decision rule [5]. Define:

- x = how long the data (or instrument) must remain secure—its security shelf-life;
- y = how long migration to quantum-safe cryptography will take;
- z = how long until a CRQC capable of breaking today's cryptography exists.

The institution has a problem if

$$x + y > z. \tag{1}$$

The reading is operational, not academic. By the time the institution finishes migrating (which takes y), data harvested today must still be protected for the remainder of its shelf-life x. If a CRQC arrives at z before that combined horizon elapses, harvested material is exposed. The only controllable terms are x (shorten by re-keying, re-issuing, or expiring data) and y (shorten by investing in crypto-agility *now*). z is set by adversary capability and is not under the institution's control; it must be treated as a distribution, not a date.

Why IT framing fails. Treated as an IT line item, migration is a discretionary project competing for budget against the next quarter's features, and Equation (1) is invisible. Treated as a balance-sheet item, the same migration is a hedge against a contingent loss whose magnitude scales with the institution's long-dated exposure. The CFO's question is not "should IT upgrade crypto?" but "what notional am I short the quantum assumption on, and at what cost can I close the position?"

3 A Simple Risk Model (Illustrative)

We price the exposure with the standard expected-loss decomposition. The framework below is what matters; *every numeric value is illustrative and is not an ACM measurement or forecast.*

Expected loss. For a single instrument or portfolio bucket,

$$\mathbb{E}[\text{loss}] \;\approx\; E \cdot \mathbb{P}(\text{break before maturity}) \cdot L, \tag{2}$$

where E is the exposure at risk (notional plus the cost of a forged transfer, failed settlement, or compromised custody record), $L \in [0, 1]$ is the loss-given-break (the fraction of E actually lost once a break occurs—rarely 1, because legal, operational, and insurance backstops recover part), and $\mathbb{P}$(break before maturity) is the probability that a CRQC arrives, and is used against this instrument, before its tenor T_{mat} elapses.

CRQC arrival as a distribution. Let Z be the (uncertain) year a CRQC becomes available, with cumulative distribution $F_Z(t) = \mathbb{P}(Z \le t)$. For an instrument maturing at T_{mat} with *no* migration, the break probability is simply

$$\mathbb{P}(\text{break before maturity}) \;=\; F_Z(T_{\text{mat}}). \tag{3}$$

F_Z is genuinely unknown and must be elicited as a subjective distribution and stress-tested, exactly as institutions already do for catastrophe and tail risk. A defensible practice is to carry several scenarios—e.g. a median CRQC horizon and a pessimistic early-arrival tail—and report $\mathbb{E}$[loss] under each rather than a single point estimate.

Migration shortens cryptographic duration. Migration does not change Z; it changes the instrument's *exposure window*. Suppose the institution completes migration to quantum-safe cryptography at time τ (its realized y, measured from today). After τ, newly issued and re-keyed material is no longer breakable by Shor, so the only exposure that remains is to material *harvested before* τ whose secrets are still live. Two regimes follow:

- **Forward secrecy (signatures, keys re-issued at τ).** Once keys rotate to post-quantum algorithms at τ, a CRQC arriving later cannot forge against the new keys. The break probability collapses from $F_Z(T_{\text{mat}})$ to $F_Z(\tau)$:

  $$\mathbb{P}(\text{break before maturity}) \;\longrightarrow\; F_Z(\tau), \qquad \tau \le T_{\text{mat}}. \tag{4}$$

  The instrument's *cryptographic duration*—the horizon over which it is short the quantum assumption—falls from T_{mat} to τ. Migrating sooner (smaller τ) is the lever, and it is the same lever as shrinking y in Equation (1).

- **Harvested confidentiality (data recorded before τ).** Ciphertext already on the wire or on-chain before τ remains exposed regardless of later migration; for that material the only mitigations are re-encryption under a quantum-safe key before harvest is useful, or reducing the secret's shelf-life x. This is why confidentiality migration is more urgent than signature migration: you cannot re-key the past.

An illustrative calculation. Consider a single 30-year bond with $E = \$100\text{m}$ at-risk notional and a realized loss-given-break $L = 0.5$ (illustrative). Suppose an elicited CRQC distribution assigns $F_Z(30\text{yr}) = 0.40$ and $F_Z(\tau) = 0.05$ for a near-term migration completing at $\tau = 3$ years (illustrative). Without migration, Equation (2) gives $\mathbb{E}[\text{loss}] \approx \$100\text{m} \times 0.40 \times 0.5 = \$20\text{m}$. With migration at $\tau = 3$, Equation (4) gives $\mathbb{E}[\text{loss}] \approx \$100\text{m} \times 0.05 \times 0.5 = \$2.5\text{m}$. The *risk reduction*, \$17.5m of expected loss on a single position, is the economic value of crypto-agility for that instrument; compared against migration cost it yields a return-on-migration the CFO can underwrite. (Numbers illustrative; substitute the institution's own exposures and elicited F_Z.)

4 Governance and Sequencing

The control objective is *crypto-agility*: the ability to change cryptographic primitives without re-architecting the systems that depend on them. An institution that can swap algorithms behind a stable interface has a small y in Equation (1); one that has hard-wired ECDSA into every signing path has a large y and a correspondingly large exposure window. Crypto-agility is the capability that makes every other mitigation in this paper cheap.

Sequencing by cryptographic duration. Not all systems carry the same x. Migrate in order of $x + y - z$ risk, which in practice means:

1. **Long-term archives and confidentiality ($x = \infty$, harvest-now-dominant).** Records that must stay confidential for decades and are exposed on the wire or on-chain *today* are highest priority, because past harvest cannot be undone. Re-encrypt under quantum-safe key establishment before more material accumulates.

2. **Custody and key management.** The roots of trust that control asset movement: long-lived signing keys, threshold-custody material, and HSM/KMS wrapping keys. A compromise here is unbounded, so these get conservative, long-horizon post-quantum protection.

3. **Settlement and signing.** Transaction authorization and settlement finality. Hybrid signatures (below) let these migrate without a flag-day cutover.

4. **Identity.** Credentials, certificates, and DID chains that authenticate counterparties; these gate everything above and should migrate in lockstep with custody and settlement.

Hybrid signatures as a hedge. During transition, a hybrid scheme attaches *both* a classical and a post-quantum signature to each artifact and accepts it if *either* verifies. The composition is secure as long as at least one underlying assumption holds, so the institution is protected against (a) a quantum break of the classical scheme and (b) an as-yet-undiscovered weakness in the newer post-quantum scheme, simultaneously [6, 8]. Operationally, hybrid removes the flag-day risk: classical and post-quantum verifiers coexist, counterparties upgrade asynchronously, and the institution never faces a moment where a single algorithm failure halts settlement. In balance-sheet terms, hybrid is a cheap option that caps downside during the window τ when migration is in flight.

What the board should see. A defensible quantum-risk posture reports, per portfolio bucket: the at-risk notional E, the instrument tenor T_{mat}, the current cryptographic duration (today's τ if migration is incomplete), the elicited F_Z scenarios, and the resulting $\mathbb{E}[\text{loss}]$ before and after planned migration. That is a quarterly risk report, not an IT status update.

5 ACM's Crypto-Agile Answer

ACM Global Tech licenses and integrates an end-to-end, post-quantum, regulatory-compliant infrastructure stack, built on Lux (blockchain, post-quantum cryptography, FHE), Hanzo (AI/data), and Zen (LLMs). The design goals below map directly to the levers in the model above.

Post-quantum from genesis (τ at issuance).
The three NIST standards—ML-KEM (FIPS 203) [1], ML-DSA (FIPS 204) [2], and SLH-DSA (FIPS 205) [3]—are deployed as native ledger precompiles activated at the genesis block, not

retrofitted [6]. For instruments issued on the platform, cryptographic duration is set to zero at issuance: there is no exposed classical key to harvest, so $F_Z(\tau)$ in Equation (4) is driven toward $F_Z(0)$. Native precompiles keep post-quantum verification economically feasible on-chain rather than prohibitively expensive in pure contract code [7].

Crypto-agility as an interface (y minimized).
Signing, key establishment, and key wrapping sit behind stable interfaces, so the underlying primitive can be swapped without changing callers—the capability that keeps migration time y small for adopting institutions [8].

Hybrid migration (option against either-side failure).
Transactions and credentials may carry both classical and post-quantum signatures and are accepted if either verifies, so existing classical workflows migrate without a cutover and stay secure as long as either assumption holds [6, 8].

Compliant platform around the cryptography.
As a Web3 Alliance (W3A, w3a.foundation) member, ACM licenses and resells this IP and co-builds with partners; the platform pairs the post-quantum core with a real-time AI/ML AML and transaction-monitoring engine and a Stripe-class payment service provider bundled with a full Banking-as-a-Service offering, so institutions adopt quantum-safe cryptography as part of a complete, regulated stack rather than as a standalone project.

These are *design goals* realized by the cited research stack; specific ACM deployment outcomes, client engagements, and production metrics are not asserted here.

6 Conclusion

For a long-dated instrument, cryptography is not infrastructure plumbing; it is an assumption baked into the instrument's value, with a duration of its own. Mosca's inequality $x + y > z$ tells the institution whether it is already short that assumption, and the expected-loss decomposition $\mathbb{E}[\text{loss}] \approx E \cdot F_Z(\cdot) \cdot L$ tells it how much. Migration is the hedge: it shortens cryptographic duration from the instrument's full tenor to the migration horizon τ, and crypto-agility is what makes τ small. ACM's contribution is to make τ effectively zero for new issuance—post-quantum from genesis, hybrid during transition—so that the quantum threat is priced, governed, and hedged like any other duration risk on the balance sheet.

References

[1] National Institute of Standards and Technology. *FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard* (ML-KEM). U.S. Department of Commerce, 2024. https://csrc.nist.gov/pubs/fips/203/final.

[2] National Institute of Standards and Technology. *FIPS 204: Module-Lattice-Based Digital Signature Standard* (ML-DSA). U.S. Department of Commerce, 2024. https://csrc.nist.gov/pubs/fips/204/final.

[3] National Institute of Standards and Technology. *FIPS 205: Stateless Hash-Based Digital Signature Standard* (SLH-DSA). U.S. Department of Commerce, 2024. https://csrc.nist.gov/pubs/fips/205/final.

[4] National Institute of Standards and Technology. *Post-Quantum Cryptography Standardization* (project overview; first standards finalized 2024). https://csrc.nist.gov/projects/post-quantum-cryptography.

[5] M. Mosca. *Cybersecurity in an Era with Quantum Computers: Will We Be Ready?* IEEE Security & Privacy, vol. 16, no. 5, pp. 38–41, 2018. States the risk inequality $x + y > z$ relating data security shelf-life (x), migration time (y), and time to a cryptographically relevant quantum computer (z).

[6] Z. Kelling. *Post-Quantum Security for Digital Securities Infrastructure*. Lux Industries, Inc. Source: `lux:post-quantum-securities/post-quantum-securities.tex`.

[7] Z. Kelling. *Post-Quantum Cryptographic Suite for EVM: ML-KEM, ML-DSA, and SLH-DSA as Native Precompiles*. Lux Industries, Inc. Source: `lux:lux-pq-crypto-suite.tex`.

[8] Z. Kelling. *Hanzo Post-Quantum Cryptography: Migration Path for AI Infrastructure*. Hanzo Industries Inc. Source: `hanzo:hanzo-pq-crypto/hanzo-pq-crypto.tex`.