

Five Years to Quantum-Safe: Meeting the 2030/2031 U.S. Federal Post-Quantum Cryptography Mandate

ACM Global Tech

Abstract

On June 23, 2026, the White House issued an executive order, *Securing the Nation against Advanced Cryptographic Attacks* [1], that compresses the U.S. federal post-quantum cryptography timeline by roughly four to five years. Computing systems supporting “high-value assets” and “high-impact systems” must move to post-quantum *key establishment* by December 31, 2030, and to quantum-safe *digital signatures* by December 31, 2031. Under the prior plan—the NSA’s 2022 CNSA 2.0 timeline [3]—National Security Systems had until 2030–2033 and most other organizations until 2035. The order also stands up a government-wide transition coordination process, directs NIST and CISA to issue guidance on a *cryptographic bill of materials* (CBOM), and binds “covered contractors” to the same deadlines through procurement rules [1, 2]. A companion order funds a national effort to build a quantum computer powerful enough to “initiate the era of quantum-enabled scientific discovery” [2]. The catalyst is a falling cost estimate for the attack itself: 2026 research describes breaking ECC-256 with roughly 30,000 physical qubits, and a Google team reports elliptic-curve discrete-log circuits at about 500,000 physical qubits—about half the team’s prior figure for breaking RSA-2048, which itself fell from a billion qubits in 2012 to roughly 20 million by 2019 [2, 5]. For regulated finance the message is no longer aspirational: crypto-agility is now a compliance deadline, not a research topic. This paper translates the mandate into the capabilities an institution must field, maps each requirement to a migration step, and describes how ACM Global Tech’s stack—NIST-standardized post-quantum algorithms as native precompiles, with hybrid migration, built on Lux, Hanzo, and Zen—is already positioned at the 2030/2031 target.

1 The Mandate, in Plain Terms

The executive order *Securing the Nation against Advanced Cryptographic Attacks* [1] sets two hard dates. By **December 31, 2030**, in-scope federal computing systems must perform key establishment—the handshake that negotiates a shared symmetric key for a session—using post-quantum schemes. By **December 31, 2031**, those systems must authenticate using quantum-safe digital signatures. The two deadlines are split deliberately: key establishment is the more urgent because a recorded handshake can be broken retroactively, whereas a signature only needs to resist forgery while it is still trusted.

Who is covered. The order reaches three populations. First, systems supporting *high-value assets*—the data and applications whose loss would cause severe harm to agency missions. Second, *high-impact systems*, the federal categorization for the most sensitive workloads. Third, and most consequential for the private sector, *covered contractors*: procurement rules bind vendors selling into the government to the same 2030/2031 deadlines and require them to maintain vulnerability

disclosure policies [1, 2]. A bank, custodian, or payments provider that serves any federal counterparty, or that aspires to, inherits the deadline by contract even though it is not itself a federal agency.

Years sooner than before. The prior baseline was the NSA’s Commercial National Security Algorithm Suite 2.0, published in 2022 [3], under which National Security Systems were to be quantum-ready over 2030–2033 and the broader ecosystem had until roughly 2035. Pulling the general deadline forward to 2030/2031 removes four to five years of slack. For an institution that had penciled migration into a 2033–2035 capital plan, the runway just shrank to under five years—and much of that window is consumed by discovery and testing, not the cryptographic swap itself.

What it means for regulated finance. The practical translation is blunt: *crypto-agility is now a compliance deadline, not a research topic*. The question a Chief Information Security Officer must answer is no longer “when might quantum matter?” but “can we evidence, to an examiner, a post-quantum key-establishment posture by 2030 and a quantum-safe signature posture by 2031, across every system in scope?” That is a program-management and controls question with a fixed due date, and it begins with knowing what cryptography the institution actually runs—which most institutions today cannot enumerate.

2 Why Now—The Threat Moved

A deadline is only as credible as the threat behind it, and the threat estimate has moved sharply in the institution’s disfavor.

Falling qubit estimates. The resource cost of the attack is collapsing. In 2012, breaking RSA-2048 was estimated to need on the order of a billion physical qubits; by 2019, Gidney and Ekerå put it at roughly 20 million noisy qubits in eight hours [5]. In March 2026, researchers described recovering an ECC-256 private key—the curve securing Bitcoin and Ethereum—using only about 30,000 physical qubits in roughly ten days, and a Google research team described quantum circuits solving the elliptic-curve discrete-logarithm problem with about 500,000 physical qubits, roughly half the same team’s earlier estimate for RSA-2048 [2]. Each revision has moved the bar lower, not higher. The trend, not any single number, is the signal: the hardware target a defender must out-run is receding toward them.

Industry already moved its own clocks. This is not only a government reading. Google and Cloudflare—operators of some of the largest cryptographic estates in existence—recently advanced their own post-quantum migration timelines to 2029 [2]. When the parties with the deepest visibility into protocol-level cryptography set internal deadlines ahead of the federal one, a 2030 mandate looks like a floor, not a ceiling.

The mathematics is settled. The vulnerability is not speculative. Shor’s algorithm [4] factors integers (breaking RSA) and computes discrete logarithms (breaking ECDSA and Diffie–Hellman) in polynomial—cubic—time on a sufficiently large quantum computer. A cryptographically relevant quantum computer therefore recovers a private key from its public key and forges signatures at will. The NIST post-quantum standards [7, 8, 9] rest instead on lattice and hash problems for which no comparable quantum advantage is known. The open variable is engineering—qubit count and error correction—and that variable is the one moving against the defender.

Harvest now, decrypt later. The most acute risk for finance needs no quantum computer today. A patient adversary records ciphertext, public keys, and signatures now, stores them cheaply, and decrypts them the day a CRQC exists. For long-lived financial records—multi-decade custody trails, 30-year bonds, archived settlement data, identity credentials with long validity—the secret recorded in 2026 may still need to be confidential in 2040. Harvest-now-decrypt-later (HNDL) collapses the comfortable distance between “the machine does not exist yet” and “my data is already exposed.”

Mosca’s inequality. Michele Mosca’s framing makes the timing arithmetic explicit [6]. Let

- x = how long the data must remain secret (its 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 whenever

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

The reading is operational. Data harvested today must stay protected for the remainder of its shelf-life x even after migration finishes y years from now; if a CRQC arrives at z before that combined horizon elapses, harvested material is exposed. A federal deadline of 2030 is, in effect, a regulatory assertion that z is close enough that $x + y$ already exceeds it for in-scope data. For a financial record with x measured in decades, even a modest y pushes $x + y$ past any plausible z —which is precisely why a 2030 deadline implies *starting now*, not in 2029. The only terms the institution controls are x (shorten by re-keying or expiring data) and y (shorten by investing in crypto-agility today); z is set by the adversary and, as the estimates above show, is shrinking.

3 What the EO Demands, Operationally

Stripped to its operational core, the order requires five capabilities. Each maps one-to-one to something an institution must be able to do.

(a) Post-quantum key establishment by 2030. Every protocol that negotiates a session key—TLS, VPN tunnels, SSH, application-layer key exchange—must use a post-quantum key-encapsulation mechanism. The NIST standard is ML-KEM (FIPS 203) [7]. This is the harder of the two deadlines to ignore, because a handshake recorded today is breakable the moment a CRQC exists, so confidentiality migration is retroactively urgent.

(b) Quantum-safe signatures by 2031. Every system that authenticates—code signing, certificate chains, document and transaction signing, firmware—must move to a NIST signature standard: ML-DSA (FIPS 204) for general use, or SLH-DSA (FIPS 205), a conservative hash-based scheme, where a different security assumption is warranted [8, 9].

(c) A cryptographic bill of materials (CBOM). The order directs NIST and CISA to issue guidance on a CBOM: an inventory of every cryptographic component, library, and module in an encryption system [1, 2]. Conceptually it is to cryptography what a software bill of materials is to dependencies—a complete, machine-readable list of what algorithms, key sizes, certificates,

and crypto libraries are in use, and where. The requirement exposes an uncomfortable fact: *most institutions cannot today enumerate their own cryptography*. A typical estate has TLS terminating in a dozen places, certificates issued by overlapping authorities, embedded crypto in third-party appliances, key material in HSMs and cloud KMS, and signing logic compiled into binaries nobody has audited in years. You cannot migrate what you cannot list, and the CBOM is the order’s mechanism for forcing the list to exist.

(d) Crypto-agility. Implicit in two deadlines two years apart, and in the standing transition-coordination process the order establishes [1], is the requirement that primitives can be *rotated*—swapped behind stable interfaces without re-architecting the systems that depend on them. An institution that has hard-wired ECDSA into every signing path faces a large y in Equation (1); one that abstracts the primitive can rotate ML-KEM in 2030 and ML-DSA in 2031 without a second migration program. Crypto-agility is also the hedge against the next transition: NIST will deprecate and replace algorithms again, and an agile institution treats that as routine.

(e) Vulnerability disclosure for contractors. Covered contractors must implement vulnerability disclosure policies as a condition of sale [1]. For a vendor selling into regulated finance or government, this folds cryptographic posture into the same disclosure and remediation discipline already expected for software vulnerabilities.

Government-wide coordination. Above these five, the order installs a transition-coordination process led by the OMB Director and the National Cyber Director, with each agency naming a point person, and directs the State Department—with NIST, DoD, DHS, the National Cyber Director, and the DNI—to engage foreign governments and industry to adopt NIST-standardized PQC [1]. The international push matters to any institution with cross-border rails: the standards an institution adopts in the U.S. are the ones it will be expected to interoperate on abroad.

4 How ACM Meets the Mandate Today

ACM Global Tech’s position is unusual: the capabilities the order demands by 2030 and 2031 are deployed in ACM’s stack now. ACM 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), and as a Web3 Alliance (W3A, w3a.foundation) member, ACM licenses this IP and partners to resell and co-build it for regulated institutions.

The three NIST algorithms, native, from genesis.

ML-KEM (FIPS 203), ML-DSA (FIPS 204), and SLH-DSA (FIPS 205) are deployed as native EVM precompiles activated at the genesis block of the Lux post-quantum ledger, not retrofitted [7, 8, 9, 10]. This is precisely the (a) key-establishment and (b) signature capability the order requires—present today, not on a 2030 roadmap. Native precompiles keep post-quantum verification economically feasible on-chain rather than prohibitively expensive in pure contract code [11].

Hybrid migration, non-disruptive by construction.

During transition, an artifact may carry both a classical and a post-quantum signature and is accepted if *either* verifies. The composition is secure as long as at least one assumption holds, so classical and post-quantum verifiers coexist and counterparties upgrade asynchronously—there

is no flag-day cutover [10, 12]. This is the mechanism that lets an institution satisfy the 2030 and 2031 deadlines without halting settlement on a cutover date.

Crypto-agility as an interface.

Signing, key establishment, and key wrapping sit behind stable interfaces, so the underlying primitive can be rotated without changing callers [12]. This is capability (d)—the property that keeps migration time y small now and at the next transition.

Threshold custody, privacy, and AML around the cryptography.

MPC threshold custody distributes signing authority so no single party can move assets; FHE supports computation over encrypted data; and ACM operates a real-time AI/ML AML and transaction-monitoring engine, all wrapped in a Stripe-class payment service provider with full Banking-as-a-Service. Institutions adopt quantum-safe cryptography as part of a complete, regulated stack rather than as a standalone project.

On the CBOM. The order’s CBOM requirement (c) is a discovery problem before it is a cryptography problem: an institution must inventory every crypto component, library, and module it runs. ACM’s platform can generate a structured cryptographic inventory for the systems it operates and integrates, and ACM can run a discovery engagement to enumerate an institution’s quantum-vulnerable cryptography and prioritize it by exposure. Where a system is outside ACM’s stack—legacy appliances, third-party software, embedded crypto—CBOM production is a scoped discovery exercise, not an automatic output.

5 A Migration Sequence for a Regulated Institution

The order’s requirements assemble into a single sequence. The outcomes below are stated as design goals and program targets, not guarantees.

1. **Inventory (CBOM).** Enumerate every cryptographic component, library, certificate, and key store across the estate—including third-party and embedded crypto. This is the order’s CBOM requirement and the precondition for everything after it: you cannot migrate what you cannot list. Target outcome: a complete, machine-readable cryptographic inventory mapped to systems and data flows.
2. **Prioritize by secrecy-lifetime and exposure (HN DL).** Rank the inventory by x (how long each secret must stay confidential) and by whether the material is already exposed to harvest—on the wire, on a public ledger, or in shared infrastructure. Long-lived confidential data that is exposed today ranks highest, because past harvest cannot be undone. This is Mosca’s $x + y > z$ applied per system.
3. **Deploy hybrid PQC for key establishment first (2030).** Move key establishment to ML-KEM under a hybrid (classical + PQ) scheme, so confidentiality migrates without a cutover and recorded handshakes stop being a retroactive liability. This satisfies the 2030 deadline and addresses the most urgent HN DL exposure first.
4. **Migrate signatures (2031).** Move authentication to ML-DSA, or SLH-DSA where a hash-based assumption is warranted, again under a hybrid scheme so classical and post-quantum verifiers coexist during the transition. This satisfies the 2031 deadline.

5. **Establish crypto-agility as an ongoing control.** Abstract primitives behind stable interfaces so future rotations are routine, fold cryptographic posture into vulnerability disclosure and examiner reporting, and keep the CBOM live rather than treating it as a one-time artifact. The target is an institution for which the *next* transition is configuration, not a program.

6 Conclusion

The executive order is a forcing function. It converts a risk that institutions could defer into a dated compliance obligation—2030 for key establishment, 2031 for signatures—backed by procurement rules that reach private-sector vendors and by a threat estimate that keeps falling. The institutions that meet it will be the ones that already treat cryptography as *agile infrastructure*: inventoried, abstracted behind stable interfaces, and rotatable on a schedule, rather than hard-wired and swapped once in a panic. ACM Global Tech’s contribution is to provide that substrate—the three NIST algorithms native from genesis, hybrid migration that avoids a cutover, crypto-agility as an interface, and a regulated platform around the cryptography—built on Lux, Hanzo, and Zen. The deadline is five years out. The work of knowing what you run, and making it rotatable, starts now.

References

- [1] The White House. *Securing the Nation against Advanced Cryptographic Attacks*. Executive Order, 2026. (Sets the December 31, 2030 deadline for post-quantum key establishment and the December 31, 2031 deadline for quantum-safe digital signatures for high-value assets, high-impact systems, and covered contractors; establishes a government-wide transition coordination process and directs NIST/CISA guidance on a cryptographic bill of materials. A companion executive order, issued the same day, directs federal support for quantum computing, including a national effort to build a quantum computer capable of initiating quantum-enabled scientific discovery.)
- [2] D. Goodin. *White House drastically shortens deadline for dropping quantum-vulnerable crypto*. Ars Technica, June 2026. (Reporting on the executive order and on the falling qubit estimates, including the March 2026 ECC-256 result at roughly 30,000 physical qubits, the Google elliptic-curve discrete-log circuits at roughly 500,000 physical qubits, and Google’s and Cloudflare’s move to 2029 migration timelines.)
- [3] National Security Agency. *Commercial National Security Algorithm Suite 2.0 (CNSA 2.0)*. 2022. (The prior timeline under which National Security Systems were to be quantum-ready over 2030–2033 and most other organizations had until roughly 2035.)
- [4] P. Shor. *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*. SIAM Journal on Computing, vol. 26, no. 5, pp. 1484–1509, 1997.
- [5] C. Gidney and M. Ekerå. *How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits*. 2019.
- [6] 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).

- [7] 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>.
- [8] 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>.
- [9] 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>.
- [10] Z. Kelling. *Post-Quantum Security for Digital Securities Infrastructure*. Lux Industries, Inc. Source: `lux:post-quantum-securities/post-quantum-securities.tex`.
- [11] 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`.
- [12] Z. Kelling. *Hanzo Post-Quantum Cryptography: Migration Path for AI Infrastructure*. Hanzo Industries Inc. Source: `hanzo:hanzo-pq-crypto/hanzo-pq-crypto.tex`.