



**FRONTIER.**

FRONTIER RESEARCH / DECISION-GRADE ANALYSIS

# When the Keys Fail: Surviving the Enterprise Security Gauntlet

ECC is not broken. Quantum attack costs are being re-estimated downward, while AI is industrializing software vulnerability discovery. A fact-grounded CIO/CISO blueprint for protecting sensitive data, identity, and operations even when today's trust assumptions fail.

## REPORT PROFILE

# Built to change a decision.

ECC is not broken. Quantum attack costs are being re-estimated downward, while AI is industrializing software vulnerability discovery. A fact-grounded CIO/CISO blueprint for protecting sensitive data, identity, and operations even when today's trust assumptions fail.

---

|        |        |
|--------|--------|
| REPORT | FR-028 |
|--------|--------|

---

|        |                  |
|--------|------------------|
| AUTHOR | Dion Hinchcliffe |
|--------|------------------|

---

|           |                 |
|-----------|-----------------|
| PUBLISHED | October 9, 2026 |
|-----------|-----------------|

---

|         |                 |
|---------|-----------------|
| PROFILE | frontier-ro-0.1 |
|---------|-----------------|

This PDF is a publication rendering of the canonical Frontier Research Object. The web report remains the live source of truth; citations, figures, methodology, and rights metadata travel with this edition.

---

## Executive brief: The keys are not broken. The timetable is.

Public-key cryptography still works. That sentence matters. As of October 9, 2026, there is no publicly validated, practical classical attack that recovers arbitrary private keys from properly implemented, mainstream 256-bit elliptic-curve systems at scale. Nor has a cryptographically relevant, fault-tolerant quantum computer been publicly demonstrated breaking such a system. Anyone claiming otherwise owes the security community a reproducible exploit, not an alarming slide deck.

But the absence of a mathematical break is no longer a sufficient enterprise security strategy.

Three clocks have started running at different speeds. The exploitation clock is collapsing as AI finds software weaknesses faster and increasingly helps construct exploits. The confidentiality clock is already running for information intercepted today that could be decrypted years later. The cryptographic-transition clock is measured in years because certificates, embedded devices, firmware roots, suppliers, applications, and identity systems cannot all be upgraded overnight. The strategic danger is not that these clocks show the same time. It is that they don't.[1][2][3]

A fourth, harder-to-quantify clock may matter: frontier AI is beginning to produce original mathematical research at a scale that would have sounded fanciful a year ago. OpenAI's October 6 release includes mathematical manuscripts and computer-checked Lean formalizations for many results. That is a genuine research signal; it is not evidence that AI has found a practical classical algorithm for breaking ECC. The degree to which future mathematical advances will translate into cryptanalysis is unknown. A security investment case should acknowledge that uncertainty rather than convert it into an invented countdown.[4]

Frontier thesis: The coming security crisis is not necessarily the sudden death of encryption. It is the growing possibility that an enterprise's assumptions about keys, trusted software, and time-to-remediation become invalid faster than its infrastructure can adapt. A cipher can remain mathematically strong while the institution around it becomes indefensible.

The enterprise mandate is to defend the data, authority, and business process even if an algorithm, implementation, certificate, endpoint, supplier, or operational assumption fails. This requires a coordinated multiyear program in post-quantum migration, cryptographic agility, device modernization, zero trust, attack-surface reduction, AI-enabled defense, resilient recovery, and data-life-cycle governance. In complex estates, the cost could be substantial. The defensible strategy is neither indiscriminate hardware replacement nor buying a box labeled "quantum safe." It is replacing the specific dependencies that prevent an organization from surviving cryptographic and exploitation failures.

## Six findings the board should understand

1. ECC's quantum risk is concrete; a classical break is not demonstrated. A September 2026 paper estimates that a hypothetical fault-tolerant trapped-ion system with 19,397 physical qubits could recover a secp256k1 discrete logarithm in roughly 25.7 days with 63% success under its stated architecture. This is a research resource model, not a functioning attack on today's hardware.[5]
2. The software attack surface is already under strain. Anthropic says Project Glasswing partners verified at least 129,000 vulnerabilities during April–July 2026, with another 5,500 verified in Anthropic's own April–October scanning; more than 33,000 verified findings had high or critical ratings across the reported efforts. These are vendor-reported program totals—not 129,000 exploited zero-days, not independent prevalence estimates, and not proof of ECC weakness.[1]

3. Breach pressure is measurable. Verizon's 2026 DBIR identifies exploitation of software vulnerabilities as the entry point in 31% of breaches in its studied dataset. Its public-sector analysis reports that only 26% of sampled CISA Known Exploited Vulnerability items were fully remediated in the measured organizations and a median of 43 days to full resolution. These numbers have different denominators and must not be combined into a single probability.[3][6]
4. Migration technology is real, but uneven. NIST standardized ML-KEM, ML-DSA, and SLH-DSA in August 2024. IETF RFC 10024 standardized three hybrid TLS 1.3 key-agreement groups in August 2026. Public-key signatures across PKI, firmware, hardware roots, OT, and legacy protocols will take much longer to transition than a modern web endpoint.[7][8][2]
5. The official planning horizon is already here. The UK NCSC targets estate-wide discovery and a plan by 2028, highest-priority migration by 2031, and broad completion by 2035. These are NCSC guidance milestones, not a worldwide prediction of when cryptography breaks.[2]
6. Replacement economics are unavoidable but heterogeneous. A software-upgradable browser is not a decade-old industrial controller with a fixed cryptographic trust anchor. Some organizations will replace a large fraction of devices; others can upgrade libraries, gateways, certificates, and firmware. The deciding variable is crypto-agility and vendor support, not the age of the device alone.[2]

---

## I. First, stop saying mathematics is “unassailable”

Cryptography depends on mathematical structure and hardness assumptions. The underlying mathematics is rigorous; a claim that a particular problem has no tractable attack is not necessarily a mathematical theorem. RSA relies on the practical difficulty of integer factorization. Conventional elliptic-curve cryptography relies on the difficulty of a discrete-logarithm problem on selected curves. Neither problem is proven classically intractable for all time.

Correctly implemented, properly parameterized ECC remains robust against public classical methods at the intended security levels. Attackers may exploit bad randomness, nonce reuse, side channels, invalid-point handling, weak curve choices, library defects, key theft, or protocol misuse. Those are attacks on implementations and assumptions, not evidence that the general mathematical problem has been solved. The distinction is not pedantry: it determines whether the remedy is patching a library, rotating compromised keys, updating a protocol, or abandoning an entire algorithm family.

A suitably capable fault-tolerant quantum computer running Shor's algorithm would threaten classical RSA, finite-field Diffie–Hellman, ECC key exchange, and ECDSA signatures. Quantum attacks do not make well-chosen AES or hash functions magically collapse in the same way. Grover's algorithm offers a generic quadratic-search advantage in an idealized model, informing preferences for adequate symmetric-key and hash security margins, but it is not a practical demonstration of instant AES-256 decryption.

Cryptographic risk therefore has at least four distinct classes:

| Failure class                                                    | What actually fails                                                                          | Security consequence                                                | Appropriate response                                                                |
|------------------------------------------------------------------|----------------------------------------------------------------------------------------------|---------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| Future quantum attack on traditional public keys                 | RSA factoring or elliptic-curve discrete logs become tractable on a capable quantum computer | Historic confidentiality; future authentication and signatures      | Standardized PQ key establishment and signatures, crypto agility, transition of PKI |
| Breakthrough classical cryptanalysis, including AI-assisted work | A previously unknown attack reduces hardness for a deployed scheme                           | Potentially abrupt loss of confidence in a cipher family            | Algorithm diversification, rapid re-keying, protocol replacement, isolation         |
| Exploitable cryptographic implementation                         | Library, RNG, side-channel, nonce, memory or protocol defect                                 | Key recovery, forgery or disclosure in affected deployments         | Secure implementation, patching, HSMs, managed keys, verified deployment            |
| Broad software and identity compromise                           | Zero-day, credential theft, token abuse, supplier compromise                                 | Attackers use legitimate access to see plaintext or command systems | Zero trust, endpoint security, segmented authority, monitoring, recovery            |

The last row is already an everyday problem. Post-quantum encryption does not protect plaintext from an attacker who has become your database administrator.

## II. How close is ECC to the edge?

The most interesting 2026 movement is not a confirmed classical decryption exploit. It is falling resource estimates for future quantum attacks.

Thomas Häner and a multi-institution group reported a September 2026 architecture-specific estimate for solving the 256-bit secp256k1 elliptic-curve discrete logarithm problem: roughly 1,450 logical qubits, 40 million Toffoli gates, and, after mapping their circuits to a proposed trapped-ion architecture, 19,397 physical qubits, 25.7 days, and 63% estimated success. These are conditional figures; error correction, connectivity, gate operations, noise assumptions, timing, and hardware engineering are all part of the claim. It cannot be extrapolated unchanged to every ECC curve or every quantum computer.[5]

Similarly, Craig Gidney's 2025 analysis reduced a modeled RSA-2048 factoring attack to under one million noisy physical qubits in less than one week, assuming a particular square-grid architecture, 0.1% gate error rate, one-microsecond surface-code cycles and other specified parameters. The older estimate was about 20 million qubits. This is a major algorithmic/resource-estimation advance—not an observation that an RSA-2048 key was factored.[9]

Google's own quantum-error-correction researchers caution that scaling from demonstrations on roughly a hundred physical qubits to useful fault-tolerant systems remains an enormous engineering challenge.[10] Quantum hardware counts are also not directly comparable across platforms: "physical," "logical," and error-corrected computational capacity are different metrics. The relevant CIO metric is not vendor-reported qubit count but demonstrated fault-tolerant resource capacity relative to audited cryptanalytic circuits.

Frontier's assessment: the plausible time horizon is a distribution, not a date. The failure mode is planning against the median quantum forecast when protected records must remain confidential long after the median. Long-lived secrets with existing intercepted copies can lose even if the organization finishes migration before the first public quantum break. That is the logic of harvest-now-decrypt-later (HNDL).

The three clocks, kept separate

| Clock                                             | Evidence as of Oct. 9, 2026                                                                                  | Confidence                                                                                    | Board interpretation                                                  |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|
| AI-accelerated software exploitation              | Thousands of verified findings; real-world incident trends and autonomous exploit demonstrations             | High that defensive discovery capacity is rising; uncertain attacker adoption and net outcome | Reduce exposure and speed remediation now                             |
| Quantum attack on current public-key systems      | Shor's known mathematics; falling conditional resource estimates; incomplete hardware                        | High eventual threat conditional on capable fault-tolerant machines; low confidence in timing | Transition by sensitivity and data lifetime, not by a vendor deadline |
| Classical/AI mathematical break of mainstream ECC | Frontier mathematical research accelerating; no publicly validated practical break of mainstream 256-bit ECC | High uncertainty; no evidence sufficient to claim imminent break                              | Maintain algorithm agility and independent contingency capacity       |

A fourth practical constraint is the migration schedule itself. Ten years sounds generous until the enterprise discovers a million certificates, embedded firmware that cannot accept larger signatures, and suppliers whose next approved maintenance window is measured in years.

## III. AI has changed vulnerability economics before it has changed cryptographic mathematics

Anthropic's October 6 disclosure is difficult to dismiss: partner organizations uncovered at least 129,000 verified software vulnerabilities in April–July 2026; Anthropic reported an additional 5,500 verified from its open-source scans during April–October; more than 33,000 of the combined verified findings have been rated high or critical. Anthropic says its disclosed count likely understates activity because it surveyed a subset of partners. These data describe a provider-organized research program with potential selection effects, duplicated remediation work and heterogeneous severity processes—not a representative census of deployed enterprise flaws.[1]

In September 2026 Anthropic also reported that GLM-5.3, developed by Zhipu AI/Z.ai, demonstrated strong end-to-end exploit construction capabilities. The vendor's evaluations support concern about proliferation of offensive capability beyond one tightly controlled model. They do not establish that all threat groups now have reliable, autonomous exploitation of arbitrary systems.[11]

Google Project Zero and DeepMind reported a real SQLite memory-safety flaw found by Big Sleep in October 2024, with same-day developer remediation. Google subsequently described how threat intelligence combined with Big Sleep helped identify a SQLite issue believed to be headed toward exploitation. That provides an important external example: AI security research can move from synthetic benchmarks to concrete code review and incident prevention.[12]

The scale numbers must be handled honestly. A verified vulnerability is not necessarily exploitable remotely. A high CVSS rating is not proof of exploitation. A zero-day is a vulnerability exploited or unknown before a patch; it is not interchangeable with a newly discovered flaw. An AI-found proof of concept does not establish ubiquitous attacker access. The strongest evidence here is increased discovery throughput and demonstrated exploit capability, not a measured universal collapse of time-to-exploit.

Nevertheless, the patching side of the equation is worrying. Verizon reports vulnerability exploitation as the initial vector in 31% of observed breaches in its 2026 dataset; a public-sector analysis describes 26% full remediation of catalogued known-exploited vulnerabilities and a 43-day median to full remediation in the measured organizations.[3][6] The latter cannot be generalized uncritically across sectors, but it illustrates a broken queueing equation: attackers only need the one exposed path still awaiting its change window.

IBM/Ponemon's 2026 study reports a global mean breach cost of \$4.99 million and AI-enabled malicious breaches in roughly one in four cases, with those incidents averaging \$6 million. These are survey-derived association and cost estimates, not proof that AI caused the incremental cost or that every enterprise faces the same loss distribution.[13] The enterprise security budget is now competing with an attacker whose marginal search and experimentation cost may be falling.

The zinger: When the attacker can ask a machine to find a bug, your patch committee cannot keep operating like a monthly book club.

---

## IV. Why frontier mathematics matters—and what it has not shown

OpenAI reported a broad collection of new mathematical manuscripts on October 6, 2026, with a GitHub release, formalization in Lean for many results, and summaries of research method and compute. The company has separately described surprising progress on long-standing mathematical problems. Some submissions still require independent scholarly scrutiny; computer-checked formalization is meaningful evidence for a formalized theorem, not a universal certification of all surrounding claims. None of the reviewed disclosures demonstrates a practical classical attack against production ECC or ML-KEM.[4]

The intellectual implication is profound but bounded. Modern cryptographic selection has always assumed sustained expert adversarial analysis. Frontier AI could expand the population of researchers able to search obscure constructions, automate proof strategies, improve quantum circuits, or discover implementation bugs. It may also accelerate defenses, formal verification, safer cryptographic libraries, and protocol testing. Treating every mathematical result as a threat to encryption is unserious. Assuming the known search space remains fixed is equally unserious.

The appropriate response is cryptographic option value: reduce the cost of replacing an algorithm and its operational dependencies without precommitting to a speculative mathematical outcome. This is why NIST's 2025 selection of HQC as a backup KEM matters. NIST's Dustin Moody explicitly framed the code-based candidate as a hedge should lattice-based ML-KEM prove vulnerable. HQC was selected for standardization; it should not be described as already a final replacement standard. [14]

---

## V. What post-quantum cryptography solves—and what it emphatically does not

NIST's August 2024 standards establish three distinct mechanisms: FIPS 203 / ML-KEM for establishing shared secrets; FIPS 204 / ML-DSA for digital signatures; and FIPS 205 / SLH-DSA, a stateless hash-based digital-signature option. These are different jobs. A KEM does not magically replace digital signatures; a signature does not conceal stored files. NIST SP 800-227, finalized in September 2025, provides additional guidance on secure KEM use.[15][16][17][18]

This year provided an operational bridge. RFC 10024, published as an IETF Proposed Standard in August 2026, defines TLS 1.3 hybrid key-agreement groups X25519MLKEM768, SecP256r1MLKEM768, and SecP384r1MLKEM1024. They combine classical elliptic-curve ephemeral Diffie–Hellman with ML-KEM. RFC 9954 explains the general hybrid construction: properly combining independent components seeks to preserve the shared secret if one component's assumptions fail. Implementations and negotiation still require validation, supported libraries and testing.[8][19]

For confidentiality at risk from future passive decryption, a well-implemented hybrid session can be materially better than classical-only key exchange if both endpoints really negotiate it. “Our CDN offers PQC” is not evidence that every mobile app, API client, service-to-service channel, VPN, backup replication link and partner integration has end-to-end quantum-resistant protection. A classical-only segment can become the weakest point.

Just as important, hybrid key establishment is not a post-quantum identity system. Classical signatures in certificates, firmware, documents, and authenticated protocols may remain exposed until PKI, trust stores and roots transition. The UK NCSC specifically calls out WebPKI, long-lived roots, TLS ecosystem dependencies and industrial-control protocols as challenging migration domains.[2]

Neither PQC nor zero trust restores confidentiality to a secret an attacker already copied in plaintext. And neither prevents a compromised administrator from issuing legitimate-looking destructive commands. Keep encryption, identity, endpoint integrity and operational recovery as mutually reinforcing—but distinct—layers.

## Capability and gap matrix

| Surface                         | Available defensive direction                                                               | Main remaining gap                                                      |
|---------------------------------|---------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| Modern HTTPS/TLS 1.3            | Test RFC 10024-compatible hybrid KEM key agreement on both ends                             | Coverage gaps, unsupported clients, middleboxes, certificate signatures |
| Application-to-application APIs | Upgrade service mesh, gateways and runtimes; verify negotiated suites                       | Library fragmentation, pinned clients, unmanaged third parties          |
| VPN/private networking          | Prefer documented, vetted vendor quantum-resistant or hybrid key establishment as available | Different tunnel protocols, interoperability, appliance firmware        |

|                        |                                                                                         |                                                                         |
|------------------------|-----------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| Data at rest           | AES-256 with sound key handling, envelope encryption, managed HSM/KMS, compartmentation | Current key-wrapping and remote access paths; plaintext compromise      |
| Authentication and PKI | Roadmap ML-DSA/SLH-DSA certificates/signing, staged trust-anchor migration              | WebPKI coordination, token formats, certificate sizes, legacy verifiers |
| Code/firmware signing  | Inventory roots and boot chains; plan PQ signature-capable firmware and validation      | Fixed ROM/bootloaders, storage and bandwidth constraints                |
| Industrial/medical/IoT | Segregation, authenticated gateways, vendor modernization and replacement plan          | Devices with no viable modern crypto or maintenance route               |
| High-value archives    | Limit replication and retention; prioritize PQ-protected transfer and rekeying          | Already captured ciphertext cannot be recalled                          |

## VI. The investment thesis: buy survivability, not a security talisman

For boards, this is a multiyear infrastructure program rather than a security-feature procurement. NCSC's deadlines are a reasonable external planning scaffold: 2028 for complete discovery and initial planning, 2031 for the highest-priority services, and 2035 for broad migration. NCSC explicitly anticipates significant investment for some large enterprises and special difficulty in systems with long equipment lifetimes.[2] These milestones are not safe-harbor dates if a particular organization holds data that must remain secret for decades.

A useful capital model separates cryptography modernization (libraries, gateways, HSM/KMS, certificate platforms), device and OT refresh, identity/zero-trust architecture, security engineering and verification, and continuity/incident-response readiness. Avoid double counting existing scheduled refresh expenditure as new quantum cost. Some spending should be accelerated; some is genuinely incremental; some replaces planned but inadequate controls.

Illustrative sensitivity, not an industry forecast: consider an estate of 100,000 endpoints and appliances. If discovery shows 8%, 15%, or 25% cannot support required security controls or cryptographic upgrades, and a fully loaded replacement/rollout averages an assumed \$900, \$1,500, or \$2,500 respectively, the corresponding device program is \$7.2 million, \$22.5 million, or \$62.5 million. These figures are a scenario exercise, not market price observations, and exclude infrastructure engineering, downtime, OT qualification, retraining, and software subscriptions. The purpose is to expose the large sensitivity to unupgradeable share, not claim every estate must replace most devices.

The other side of the ledger is also substantial: IBM's measured 2026 global mean breach cost of \$4.99 million is not an expected loss for your enterprise. Probability-weighted loss requires sector-specific exposure, impact scenarios, recurrence, and control efficacy. A proper board investment case compares avoided downtime, confidential-data lifetime risk, reduced attack pathways, compliance obligations, and infrastructure simplification—not a lazy multiplication of a headline breach average.[20]

Frontier position: The expensive failure is not spending too much on modern infrastructure. It is spending billions on digital transformation while leaving the organization dependent on a cryptographic trust anchor nobody can replace.

---

## VII. The enterprise defense program: nine workstreams with enforceable gates

### 1. Build the cryptographic map, not another spreadsheet inventory

Within 90 days, name a CIO/CISO joint program owner and accountable business sponsors for every tier-zero identity and trust service. Establish a cryptographic bill of materials (CBOM) and dependency graph: where RSA, ECDH, ECDSA, EdDSA and other quantum-vulnerable public-key algorithms are used; what keys they protect; who owns the trust anchor; and which devices, applications, services, third parties and records depend on it. Include TLS termination, Java and OpenSSL runtimes, SSH, S/MIME, VPN, federated identity, OAuth/JWT signatures, hardware certificates, signing pipelines, update mechanisms, backups, HSMs, OT, and mobile devices.

For each dependency record: protocol/library and version; key algorithm/size; certificate issuer and expiry; signer/verifier pair; hardware root; supplier; upgrade path; data sensitivity and useful secrecy lifetime; current negotiated algorithms; crypto-agility rating; and business criticality. Discover with configuration analysis, software composition inventories, certificate scans, packet/handshake telemetry and vendor attestations. Do not assume passive network discovery reveals crypto in encrypted application payloads.

Gate: 100% of crown-jewel data flows, tier-zero identity, signing roots and internet-facing cryptographic terminations must have named owners, known algorithms and a documented migration decision. Whole-estate inventory coverage should be measured and improved, not declared complete by fiat. NIST NCCoE's PQC migration work provides supporting discovery and testing concepts.[21]

## 2. Stop creating tomorrow's decryptable archive

Classify data by confidentiality lifetime, not just regulatory label. A payroll record with a short retention period differs from a patient history, defense design, intellectual property portfolio, industrial recipe or private signing key whose exposure has long-tail impact. Classify three dimensions: how long disclosure would harm, whether ciphertext travels through interceptable channels, and whether an attacker may already have collected it.

Prioritize modern TLS 1.3 with independently verified hybrid key agreement on high-value external and service-to-service flows. Negotiate RFC 10024-approved groups where supported; record actual negotiated suites. Verify VPNs, edge, database replication, messaging and cross-cloud links separately. Do not conflate edge-to-user protection with end-to-end protection through origin and internal hops. Require downgrade/fallback logging and clearly defined exceptions; disable obsolete protocols under a staged compatibility program.

For files at rest, use strong symmetric encryption such as AES-256 and controlled envelope encryption with rotation, compartmentalized data keys, separation of duties and HSM/KMS where warranted. Understand where key wrapping, backup export and recovery depend on RSA/ECC today. Key rotation does not make ciphertext already copied by an attacker unreadable; neither does later migrating the server to PQC. Data minimization and deletion policies reduce how much future compromise is possible, but deleting your own copies cannot erase hostile archives.

Gate: priority data-flow owners demonstrate quantum-resistant-or-hybrid key establishment where supported; remaining classical-only flows have compensating isolation, explicit risk acceptance and a funded migration date. Do not advertise "quantum safe" for authentication until its signatures and trust path also qualify.[8][2]

### 3. Rebuild identity around verified devices and constrained authority

Implement zero-trust policy enforcement for workforce, privileged, workload and agent identities. NIST SP 800-207 rejects implicit access based on network location; SP 800-207A extends granular identity-based access concepts into cloud-native and multicloud services.[22][23] Use phishing-resistant MFA, preferably hardware-bound credentials and modern managed-device posture; enforce least privilege, just-in-time privileged access, segmented administration, risk-based session limits and rapid token revocation. For workload identities, require short-lived credentials, rotation, signed deployment provenance and service-level authorization, not shared static keys.

Crucial caveat: many currently deployed FIDO2/passkey, device-attestation and token-signing ecosystems themselves rely on traditional elliptic-curve signatures. They remain powerful defenses against today's phishing and credential theft, but should not be described as post-quantum authentication merely because a passkey is hardware-backed. Track identity-provider and authenticator vendor plans for PQ signature migration; maintain token lifetime and revocation controls for interim exposure.

Gate: every tier-zero administrator and critical machine identity has phishing-resistant authentication, constrained authority, trusted device signals and tested emergency revocation; high-impact agent permissions have independent approval/containment controls. Quarterly tests must show that a compromised identity cannot laterally reach unrelated crown-jewel systems.

### 4. Modernize the endpoint and device trust base selectively but decisively

A large part of the spending will be in endpoints, servers, networking, OT and embedded devices that cannot support secure updates or modern enforcement. Build a replace, upgrade, isolate, retire decision tree. A device is a replacement candidate if its vendor no longer patches exploitable vulnerabilities; it lacks verified secure boot or a viable attestation path for its risk class; its cryptographic implementation cannot be upgraded; its key material is extractable; or it cannot support required access controls and logging. Age alone is not a sufficient criterion.

Baseline full-disk encryption; secure/measured boot where supported; hardware-backed keys; centrally managed patching; EDR/XDR; application control for critical workloads; mobile-device or unified endpoint management; local privilege restriction; exploit mitigations; and protected administrative paths. Consider memory-safe replacement for exposed services and libraries as a long-term engineering objective, not a blanket instant rewrite. Segment unpatchable OT, put protocol-aware gateways at trust boundaries and prepare vendor-approved maintenance windows. Do not push unaudited cryptographic firmware to safety-critical systems.

Gate: every critical asset can either be updated and attested, be controlled behind a demonstrably effective boundary, or have a funded retirement/replacement plan. Unknown support status is a red risk, not a default acceptance.

## 5. Turn vulnerability response into an evidence-based continuous operation

Maintain externally exposed-asset discovery, software and dependency inventories, and a validated vulnerability pipeline. Combine known exploitation signals, CISA's KEV catalog, exploit maturity, exposure, asset value and mission impact—rather than sorting a flat backlog by CVSS alone.[24] Define internal service-level objectives by risk class: for example, 24–72 hours to contain actively exploited, reachable crown-jewel flaws; 7 days for other verified internet-facing critical exposures; and 30 days for lower-exposure critical flaws, with explicit exception governance. These are Frontier proposed operating targets, not statistics or universal regulator mandates. When patching is impossible, require service isolation, virtual patching where verifiable, feature disablement or credential restriction, plus adversarial validation.

Apply AI-assisted code analysis and attack-path testing to your own estate under strict authorization, reproducible validation and vulnerability disclosure controls. Measure confirmed actionable fixes per reviewer hour, median time to validated containment, exploitable exposures per critical service, and false-positive burden. A torrent of hallucinated AI bug reports can consume the same engineers needed to fix real flaws; require minimal reproduction and maintainer-friendly evidence before escalation. NIST's Secure Software Development Framework provides a stable program-level reference.[25]

Gate: exposed exploitable paths are contained on emergency timescales, not merely entered into the patch queue. Run weekly regression tests of closure and continuous verification for assets exposed to the internet or high-value identity systems.

## 6. Engineer algorithm agility and two independent recovery paths

Adopt cryptographic interfaces that can change algorithms without replacing the application or entire network. Remove hand-coded crypto from product teams. Use vetted implementations and well-defined APIs for KEMs, signatures, key management and format negotiation. Where practicable, separate data encryption from key wrapping and application logic. Test migration rollback, mixed-mode operation and signature-format changes with production-like certificates, message sizes and traffic patterns.

Build a failure drill: pretend an enterprise-approved classical signature algorithm is suddenly disallowed; then pretend a preferred lattice KEM has a serious implementation defect. Can the organization identify affected services, deploy alternative algorithms, revoke/rotate credentials, limit exposure and restore trust before business interruption becomes unmanageable? Standards diversity helps here: NIST picked code-based HQC as a backup to lattice-based ML-KEM, explicitly to retain options under cryptanalytic surprise. Do not deploy unstandardized algorithms merely to tick a “diversity” box.[14]

Gate: complete at least one production-representative migration-and-recovery test per major trust domain. No new crown-jewel application is accepted if the algorithm and signer are immovably welded into its data format.

## 7. Reconstruct software, certificate and firmware trust over time

Inventory certificate authorities, intermediates, root stores, private code-signing keys, software update channels, bootloader ROMs, TPM/secure-element dependencies, long-lived device certificates, document signatures and notarization processes. Identify signatures that need decades of validation. Plan hybrid or parallel signing and trust-anchor renewal only using standards and product implementations appropriate to the particular ecosystem; do not assume that adding a PQ signature to a file makes an old boot ROM able to verify it.

The priority problem is integrity during transition. Old signatures may remain valid under legacy rules while new root mechanisms roll out; compromised signing authorities may also need rapid revocation and rebuilding. WebPKI is a distributed coordination challenge involving browsers, certificate authorities, root programs, logs and revocation services. OT/ICS protocols may be even harder because secure key management was missing in the first place. NCSC explicitly names these bottlenecks.[2]

Gate: signed artifacts and trust anchors have a documented verification lifetime, owner, migration standard, rollback procedure and vendor-supported recovery path. Priority firmware and code-signing infrastructure receives table-top compromise exercises annually.

## 8. Make resilience independent of perfect prevention

Use immutable, isolated and regularly restored backups, segmented recovery infrastructure, offline or logically separated administrative credentials, and clean-room rebuild paths. Preserve alternate communication channels and privileged break-glass procedures whose security does not depend on the same identity infrastructure as the production domain being restored. Encrypt backups; protect both the ciphertext and the keys; test who can invoke restoration if your normal certificate authority or SSO environment becomes untrusted.

Assume attackers can sometimes decrypt traffic, borrow identities or exploit zero-days. Design mission processes to survive those events through least privilege, separation of duties, transaction limits, anomaly detection, out-of-band approvals, fraud controls, integrity checking and independent audit. For sensitive intellectual property, limit unnecessary plaintext distribution, export and permanent retention.

Gate: twice-yearly exercises demonstrate restoration of priority services under simultaneous identity compromise, network disruption and unavailable primary keys/HSM. Recovery objectives are set from business tolerances and measured in live exercises, not copied from policy.

## 9. Turn supply-chain promises into enforceable contracts

Require suppliers to provide a dated, version-specific crypto inventory and PQ migration plan; supported firmware paths; HSM/key-lifecycle detail; test evidence for hybrid sessions; supported algorithms and downgrade behavior; signing and certificate strategy; lifecycle/end-of-support commitments; vulnerability-disclosure process; and data-location/retention controls. Make evidence—not a checkbox marked “quantum-ready”—the procurement criterion.

Require strategic cloud, identity, network, endpoint, payment, medical and OT suppliers to disclose the weakest classical-only hop in a protected transaction. Add contractual notification obligations for critical cryptographic and identity vulnerabilities; rights to test interoperability; exportability for keys/metadata where legally and technically appropriate; and exit/escrow contingencies for devices whose vendors disappear.

Gate: no long-lived strategic contract renews without a credible migration and vulnerability-response roadmap, explicit fallback obligations and a stated cost allocation for required upgrades.

# VIII. Program sequencing: do the irreversible work last, and the preparatory work now

| Horizon       | Non-negotiable outputs                                                                                                                      | Accountable lead                   |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| First 30 days | Joint CIO/CISO mandate; crown-jewel and HNDL exposure triage; executive risk register; freeze on new obsolete crypto for greenfield systems | CIO, CISO, enterprise architecture |

|             |                                                                                                                                                                                                               |                                                     |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
| Days 31–90  | Initial CBOM; prioritized classical-only flow list; RFC 10024 hybrid pilots; EOL device list; supplier questionnaire; KEV response baseline; first restore/identity drill                                     | Cryptography lead, infrastructure, SOC, procurement |
| Months 3–12 | Hybrid protection on selected high-value flows; measurable internet-facing exposure reduction; zero-trust priority enforcement; tier-zero trust-map; asset refresh capex model; first algorithm-swap exercise | CIO/CISO program board                              |
| 2027–2028   | Full dependency discovery; standard-aligned PQ deployments where interoperable; funded PKI/OT modernization; initial long-lived-data migration; enterprise plan meeting NCSC 2028 milestone                   | Architecture, finance, major vendors                |
| 2029–2031   | Highest-risk signatures, identity chains, trust anchors and long-lived devices migrated as implementations mature; annual crypto-break exercises; NCSC 2031 priority milestone                                | Business service owners                             |
| 2032–2035   | Complete remaining broad transition under applicable standards; resolve exceptions, decommission legacy roots and classical-only critical dependencies; maintain crypto-agility                               | Enterprise risk and architecture                    |

The timeline does not mean delaying long-lived data protection until 2028. A file intercepted in 2026 may already be at risk if its useful confidentiality life exceeds the eventual cryptanalytic capability horizon. Equally, it is reckless to rush immature certificates or unverified firmware into safety-critical production because a date on a PowerPoint says “PQC complete.” Sequencing must follow threat exposure and verified interoperability.[2][8]

## IX. Decision cockpit: twelve measures the board should actually review

| Metric | Definition and operational test | Target logic |
|--------|---------------------------------|--------------|
|--------|---------------------------------|--------------|

|                                    |                                                                                                       |                                                                                |
|------------------------------------|-------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
| Cryptographic discovery coverage   | Owned key/protocol dependencies inventoried ÷ estimated in-scope dependencies                         | 100% for crown jewels; improve measured estate coverage quarterly              |
| Confidentiality-lifetime exposure  | High-value datasets whose secrecy horizon outlasts assessed transition window                         | Decreasing prioritized count; document impossible-to-recall intercepted copies |
| Hybrid transport coverage          | Verified RFC 10024-style hybrid key exchange on designated sensitive flows ÷ flows eligible to deploy | Track by actual negotiation and service class, not vendor brochure             |
| Classical-only identity dependency | Tier-zero chains still relying solely on quantum-vulnerable public-key signatures                     | Accountable migration/containment for each dependency                          |
| Unupgradeable critical devices     | Number and value-weighted risk of exposed assets lacking supported patch/crypto path                  | All have dated upgrade, isolation or retirement plan                           |
| KEV containment lag                | Time from actionable verified exploitation signal to tested containment                               | Critical reachable exposures measured in hours/days, not quarters              |
| Verified critical exposure         | Number of reachable known-exploitable paths on crown-jewel services                                   | Near-zero persistent exposure; exceptions expire automatically                 |
| Algorithm-change recovery time     | Hours to swap key establishment or signature mechanisms in a rehearsed trust domain                   | Reduce with drills; no undocumented or irreversible critical dependency        |
| Privilege containment coverage     | Privileged sessions under device-attested, time-limited, monitored policy                             | 100% tier-zero, then expand                                                    |
| Supplier PQ readiness              | Critical suppliers delivering current implementation evidence, timeline and test results              | No critical unowned supplier gap                                               |
| Integrity/recovery test success    | Restores succeeding under assumed IDP/PKI compromise in timed exercises                               | Meets business-defined RTO/RPO in actual test                                  |
| Economic runway                    | Funded migration requirements versus uncovered capex/opex by year                                     | Fully financed priority wave; transparent future-year commitments              |

Don't merge the twelve measures into a single red/amber/green "quantum readiness score." That would hide precisely the failure modes the program is intended to reveal.

---

## X. Scenario planning: three adversaries, three very different crises

Scenario A — the AI exploit acceleration shock (plausible now). A commodity model materially improves end-to-end exploit chaining against common enterprise software. Attackers seize an internet-facing identity connector and pivot using valid tokens. ECC remains completely unbroken. The winning investments are exposure reduction, verified patch/containment SLAs, workload isolation, constrained privileges, detection and independent recovery. PQC does not address the initial compromise.[11][3]

Scenario B — the quantum timetable compresses (uncertain timing, serious future consequence). Demonstrated fault-tolerant computing and efficient cryptanalytic circuits arrive materially faster than the organization's migration program. Captured classical-only sessions containing long-lived secrets become vulnerable to retrospective decryption; public-key authentication may require accelerated replacement. Hybrid confidentiality already deployed protects appropriately negotiated flows if its PQ component holds. Long-lived trust roots, device certificates and legacy signatures become the expensive transition bottleneck.[5][8][2]

Scenario C — a new classical mathematical or implementation failure hits a preferred primitive (low evidentiary basis for a broad mathematical break today, nonzero contingency value). Independent researchers find a novel attack against a deployed public-key method or implementation class, perhaps with AI assistance. Emergency response depends on algorithm agility, independent fallback implementations, cryptographic diversity, short key and certificate life cycles, and an ability to revoke trust. No responsible CISO should claim probability or date from today's mathematical AI headlines.[4][14]

In every scenario, assume a second failure: an operationally important endpoint is already compromised. That single addition prevents the board from mistaking new encryption algorithms for an all-purpose security strategy.

---

## XI. What CIOs and CISOs should authorize at the next board meeting

Decision 1: Establish a jointly governed cryptographic and cyber-resilience transition program. Fund a 90-day dependency discovery and high-value data-lifetime assessment, with a standing program owner, change-control authority, quarterly audit, and a multiyear costed backlog. Require a board-visible distinction among threats demonstrated today, mathematically established but hardware-dependent, and speculative. Do not finance panic on a falsely precise “Q-day.”

Decision 2: Accelerate the controls that deliver value under every plausible future. Upgrade supported software and high-risk devices, contain known-exploited attack paths, adopt strongly enforced zero trust, practice recovery without normal identity services, and protect long-lived sensitive transport with standardized hybrid PQ key establishment where interoperable. These controls pay dividends even if ECC remains classically sound for decades.

Decision 3: Make crypto-agility and supplier modernity procurement gates. Stop buying devices, applications and services whose algorithms, certificates, trusted boot chains and runtime policy cannot be updated. Require verified migration and test artifacts before anyone marks a system “quantum safe.” Allocate actual refresh funding to unupgradeable dependencies, especially OT, PKI and long-lived hardware.

The priority is not certainty that encryption will fail. It is confidence that the organization can keep operating when a trust assumption does.

The final Frontier judgment: The mathematics may endure. Your legacy security architecture may not. Stop defending the permanence of the lock. Start defending the value of what is behind the door.

---

## Methodology and evidence limitations

This report is an October 9, 2026 analyst synthesis of primary standards (NIST and IETF), government migration guidance (NCSC and NSA), quantum cryptanalytic circuit/resource estimates, vendor-reported AI security research, a primary mathematical-AI disclosure, independent vulnerability/incident telemetry and enterprise resilience standards. Facts, research models, and Frontier recommendations are distinguished in the narrative. It is not a penetration test, a quantitative probability forecast of quantum capability, proof of a classical ECC break, or independent replication of vendor performance measurements.

The 2026 AI vulnerability counts have different coverage and observation windows than 2026 DBIR incidents; they cannot be converted into a universal zero-day rate. The quantum qubit/hour estimates are conditional architectures, not present-day deployed capabilities. Hardware and financial examples are labeled illustrations rather than observed market averages. NIST IR 8547 remains an initial public draft in the cited catalogue; NCSC's 2035 date is a migration target, not a prediction. Product conformance must be tested against actual deployed versions. The report is designed for expert technical review before publication.

---

## Sources and further technical reading

Primary sources are linked below; inline reference keys resolve to the managed Frontier CMS source registry.

1. NIST, FIPS 203: ML-KEM, August 2024.
2. NIST, FIPS 204: ML-DSA, August 2024.
3. NIST, FIPS 205: SLH-DSA, August 2024.
4. NIST, PQC publications and IR 8547 status.
5. NIST, Selection of HQC as a second KEM track, March 2025.
6. NIST, SP 800-227: KEM recommendations, September 2025.
7. NIST, IR 8547 draft transition plan, November 2024.
8. NCSC, PQC migration timelines, March 2025.
9. NCSC, PQC migration workshop, 2026.
10. IETF, RFC 10024: Hybrid TLS 1.3 groups, August 2026.
11. IETF, RFC 9954: Hybrid key exchange construction, July 2026.
12. Häner et al., 256-bit ECC quantum circuit/resource analysis, September 2026, preprint.
13. Craig Gidney, RSA-2048 quantum resource estimate, 2025.
14. Google Quantum AI, Quantum error correction engineering context.
15. OpenAI, Sharing AI progress in mathematics, October 2026.
16. Anthropic, Cyber Verification Program / Glasswing results, October 6, 2026.
17. Anthropic, GLM-5.3 cyber capability evaluation, September 2026.
18. Anthropic, Open-source vulnerability scanning program, October 8, 2026.
19. Google Project Zero, Big Sleep SQLite case, 2024.
20. Verizon, 2026 DBIR highlights.
21. Verizon, 2026 DBIR public-sector vulnerability/remediation findings.
22. IBM/Ponemon, 2026 Cost of a Data Breach.

23. IBM, 2026 data-breach cost results, July 2026.
24. NIST, SP 800-207: Zero Trust Architecture.
25. NIST, SP 800-207A: Zero Trust for cloud-native services.
26. NIST NCCoE, PQC migration research project.
27. NIST, SP 800-218: Secure Software Development Framework.
28. CISA, Known Exploited Vulnerabilities Catalog.
29. NSA, CSfC post-quantum guidance addendum, draft.
30. NIST, Cybersecurity Framework 2.0.

## SOURCES / CITATIONS

# Trace the argument.

Every source below is linked. Source numbering follows first appearance in the report.

- 01 Anthropic — Expanding the Cyber Verification Program, October 6, 2026**  
<https://www.anthropic.com/news/cyber-verification-program>
- 02 UK NCSC — Timelines for migration to post-quantum cryptography, March 2025**  
<https://www.ncsc.gov.uk/guidance/pgc-migration-timelines>
- 03 Verizon — 2026 Data Breach Investigations Report key findings**  
<https://www.verizon.com/about/news/breach-industry-wide-dbir-finds>
- 04 OpenAI — Sharing AI progress in mathematics, October 6, 2026**  
<https://openai.com/index/sharing-ai-progress-in-mathematics/>
- 05 Häner et al. — Quantum resource estimate for 256-bit elliptic-curve discrete logs, September 2026**  
<https://arxiv.org/abs/2609.05625>
- 06 Verizon — 2026 DBIR public-sector analysis, remediation measurements**  
<https://www.verizon.com/business/resources/reports/2026-dbir-public-sector-snapshot.pdf>
- 07 NIST PQC publications and transition guidance**  
<https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/publications>
- 08 IETF RFC 10024 — Post-Quantum Traditional Hybrid Key Agreement for TLS 1.3, August 2026**  
<https://www.rfc-editor.org/info/rfc10024>
- 09 Craig Gidney / Google Research — RSA-2048 quantum factoring resource estimate, 2025**  
<https://research.google/pubs/how-to-factor-2048-bit-rsa-integers-with-less-than-a-million-noisy-qubits/>
- 10 Google Quantum AI — Making quantum error correction work**  
<https://research.google/blog/making-quantum-error-correction-work/>
- 11 Anthropic — GLM-5.3 and spread of advanced cyber capabilities, September 29, 2026**  
<https://www.anthropic.com/research/glm-5-3-and-the-spread-of-advanced-cyber-capabilities>
- 12 Google Project Zero — Big Sleep finds real-world SQLite vulnerability**  
<https://projectzero.google/2024/10/from-naptime-to-big-sleep.html>
- 13 IBM — 2026 global breach costs and AI-assisted breaches, July 2026**  
<https://www.newswire.ca/news-releases/ibm-study-one-in-four-malicious-breaches-are-ai-enabled-costing-companies-6-million-on-average-878451425.html>
- 14 NIST selects HQC as backup KEM, March 2025**  
<https://www.nist.gov/news-events/news/2025/03/nist-selects-hqc-fifth-algorithm-post-quantum-encryption>
- 15 NIST FIPS 203 — ML-KEM, final August 2024**  
<https://csrc.nist.gov/pubs/fips/203/final>
- 16 NIST FIPS 204 — ML-DSA, final August 2024**  
<https://csrc.nist.gov/pubs/fips/204/final>

- 17 **NIST FIPS 205 — SLH-DSA, final August 2024**  
<https://csrc.nist.gov/pubs/fips/205/final>
- 18 **NIST SP 800-227 — Recommendations for Key-Encapsulation Mechanisms, final September 2025**  
<https://csrc.nist.gov/pubs/sp/800/227/final>
- 19 **IETF RFC 9954 — Hybrid Key Exchange in TLS 1.3, July 2026**  
<https://www.rfc-editor.org/info/rfc9954>
- 20 **IBM — 2026 Cost of a Data Breach report**  
<https://www.ibm.com/reports/data-breach>
- 21 **NIST NCCoE — Migration to Post-Quantum Cryptography project**  
<https://www.nccoe.nist.gov/projects/migration-post-quantum-cryptography>
- 22 **NIST SP 800-207 — Zero Trust Architecture**  
<https://csrc.nist.gov/pubs/sp/800/207/final>
- 23 **NIST SP 800-207A — Zero Trust for cloud-native and multicloud**  
<https://csrc.nist.gov/pubs/sp/800/207/a/final>
- 24 **CISA — Known Exploited Vulnerabilities Catalog**  
<https://www.cisa.gov/known-exploited-vulnerabilities-catalog>
- 25 **NIST SP 800-218 — Secure Software Development Framework**  
<https://csrc.nist.gov/pubs/sp/800/218/final>
- 26 **NIST IR 8547 — Transition to Post-Quantum Cryptography Standards, initial public draft**  
<https://csrc.nist.gov/pubs/ir/8547/ipd>
- 27 **UK NCSC — PQC migration workshop report, August 2026**  
<https://www.ncsc.gov.uk/blogs/post-quantum-cryptography-pqc-migration-workshop-report>
- 28 **Anthropic — Opt-in vulnerability finding for open source, October 8, 2026**  
<https://www.anthropic.com/research/launching-opt-in-vuln-finding-service-for-open-source>
- 29 **NSA — CSfC Post Quantum Cryptography Guidance Addendum (draft), CNSA 2.0 timeline**  
[https://www.nsa.gov/Portals/75/documents/resources/everyone/csfc/capability-packages/CSfC%20Post%20Quantum%20Cryptography%20Guidance%20Addendum%201\\_0%20Draft%20\\_5.pdf?ver=wCGPoDQXcJKEWTgbH8xsRA%3D%3D](https://www.nsa.gov/Portals/75/documents/resources/everyone/csfc/capability-packages/CSfC%20Post%20Quantum%20Cryptography%20Guidance%20Addendum%201_0%20Draft%20_5.pdf?ver=wCGPoDQXcJKEWTgbH8xsRA%3D%3D)
- 30 **NIST Cybersecurity Framework 2.0**  
<https://www.nist.gov/cyberframework>

**FRONTIER.**

TAKE THE NEXT STEP

# Push this further.

Research is useful when it changes a decision. Bring us the question, the constraint, or the argument you want tested.

## RESEARCH / BRIEFINGS

**[inquiries@thisisfrontier.ai](mailto:inquiries@thisisfrontier.ai)**

Briefings, advisory work, executive discussions, and research inquiries.

## REPRINTS / LICENSING

**[reprints@thisisfrontier.ai](mailto:reprints@thisisfrontier.ai)**

Reprints, enterprise distribution, excerpts, and licensing.

## FRONTIER

**[www.thisisfrontier.ai](http://www.thisisfrontier.ai)**

Research, Data Insights, News, and the live Frontier corpus.

All rights reserved. Contact Frontier for redistribution or reprint rights.

SEE FURTHER. MOVE SOONER.