JUNE 2026
APPLIED QUANTUM

THE APPLIED QUANTUM PQC MIGRATION FRAMEWORK

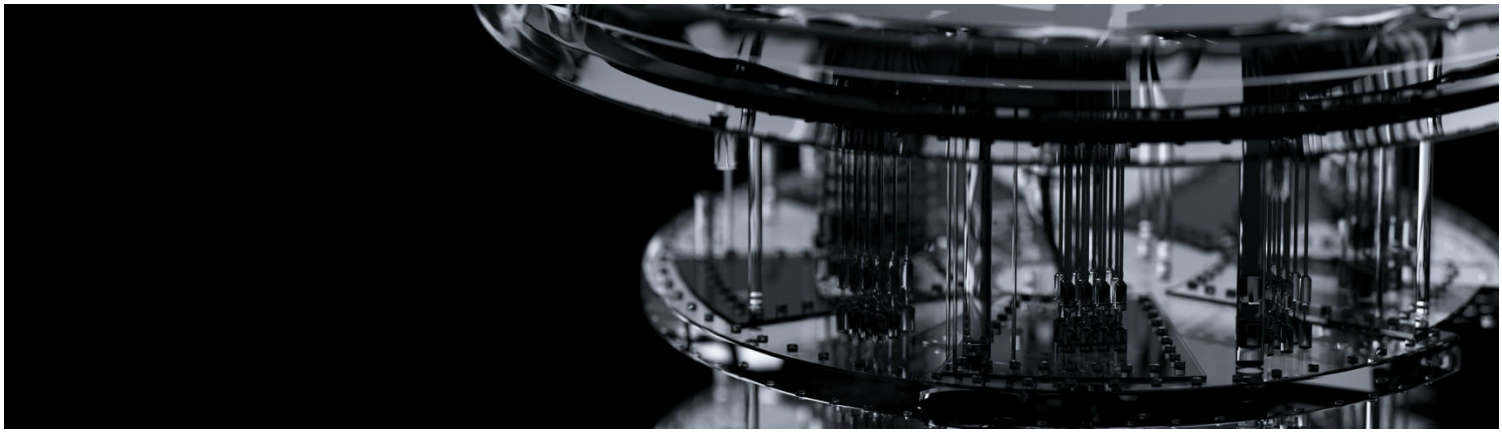

From Discovery to Deployment *- The comprehensive, phase-by-phase guide to migrating enterprise cryptography before quantum computers arrive*

Version 2.1 — June 2026

Marin Ivezic

CEO, Applied Quantum

Author, PostQuantum.com

PQCFramework.com | PQCMigrationBrief.com | PostQuantum.com | SecureQuantum.com | AppliedQuantum.com | QuantumReady.com

"Start while it's a project, before it's a crisis."

COPYRIGHT AND LICENSE

© 2026 Marin Ivezic / Applied Quantum. This work is licensed under Creative Commons Attribution 4.0 International (CC BY 4.0). You are free to share and adapt this material for any purpose, including commercial use, provided you give appropriate credit to Marin Ivezic and Applied Quantum, link to the license, and indicate any changes made.

License details: https://creativecommons.org/licenses/by/4.0/

DISCLAIMER

This framework is provided as professional guidance based on the author's and Applied Quantum's practitioner experience and publicly available standards, regulations, and industry research. It does not constitute legal, regulatory, financial, or compliance advice. Organizations should consult qualified legal counsel, auditors, and regulatory authorities for guidance specific to their jurisdiction, sector, and circumstances.

The regulatory timelines, vendor capabilities, algorithm parameters, and tool categories referenced in this document reflect the state of PQC as of June 2026. These references may become outdated quickly. Readers should verify current status against primary sources before making implementation decisions.

References to specific vendors, products, or tools are for illustrative purposes and do not constitute endorsement.

NIST IR 8547, referenced throughout for deprecation and disallowance timelines, was an Initial Public Draft as of November 2024 with final publication expected in 2026. Federal agencies and their contractors should reference the final version when available.

ABOUT THIS DOCUMENT

Document type	Enterprise methodology and practitioner framework
Intended audience	CISOs, security architects, cryptographic engineers, program managers, risk and compliance officers, vendor/procurement teams
Assumed knowledge	Familiarity with quantum threat fundamentals (HNDL, TNFL, CRQC timelines, NIST PQC standards)
Scope	Industry-agnostic core methodology with integrated operational security (SOC and GRC implementation). Sector-specific adaptations included as summary notes; dedicated sector annexes published separately.

VERSION HISTORY

Version	Date	Changes
1.0	June 2025	Initial publication. 8-phase lifecycle, cross-cutting sections, 5 appendices
1.1	Mar 2026	Reviewed and updated. 40+ updates including: NIST IR 8547 timeline, ML-DSA sizes, CNSA 2.0 milestones, LMS/XMSS status, SSH/TLS details, hybrid jurisdictional nuance. Added: library readiness, FN-DSA/HQC, Confidential Computing, legal risk, Q-by-Q Year 1 plan, dependency chains.
2.0	**Jun 2026**	Major version. Three methodological changes: (1) two-track migration model separating key exchange (HNDL) and signature/PKI (TNFL) as parallel tracks; (2) PKI architecture fork with definitive position on Merkle Tree Certificates for public Web PKI; (3) FIPS 140-3 validation gap as hard deployment constraint with environment classification. New Program Foundation sections: SOC Implementation (five detection use cases, four incident response playbooks, five tabletop exercises, three-horizon quantum CTI model), GRC Implementation (17-indicator cascading KRI framework, risk appetite statements, regulatory intelligence process, audit and assurance, GRC-SOC handoff). Expanded Program Foundations: Governance (program leadership, board oversight, three lines of defense), Crypto-Agility (five-dimensional operational discipline with four-year implementation roadmap), Skills & Team Structure (team sizing, build/borrow/buy framework, crypto champion program). Also added: cost estimation methodology, NIST CSWP 39 crypto-agility alignment, 2026–2030 regulatory deadline convergence analysis, multinational hybrid navigation framework, deployment ecosystem status data, algorithm pipeline update (9 additional signature candidates, FN-DSA/HQC status, CNSA 2.0 requirements), objection-handling appendix (Appendix F). All sector extensions updated to v2.0.
2.1	**June 2026**	Targeted update. Activity 3.3 sequencing harmonized with the two-track model; explicit position on hybrid/composite signatures; algorithm-specific vulnerability weighting added to

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

		Phase 3 risk scoring (ECC/RSA quantum resource analysis); SP 800-208 hash-based signatures foregrounded as the deploy-now component of Track B. New: Activity 2.5 (Secure the CBOM and Program Artifacts), Migration Verification & Program Closure section, identity and authentication stack in Track B scope. Added: reference program economics (0.2c), confidentiality-horizon method (1.1), evidence dossier as litigation defense, M&A due diligence, procurement model-language references, additional common failures and benefit arguments. Web PQC adoption datapoint corrected to F5 Labs mid-2025 attribution. Companion book (Quantum Ready) cross-references added. Also added: data-at-rest decision framework (5.6); AI-assisted migration position (5.7) with a Phase 1 tool category; counterparty coordination (7.6); cloud shared responsibility and SaaS (7.7); Accelerated Execution Profile (4.7); risk-weighted coverage KPI; algorithm sovereignty and standards fragmentation; crisis communications and industry information sharing (GRC); skills matrix; role-based reading paths and right-sizing profiles; Appendix G (framework crosswalk); Appendix H (protocol coverage matrix).

HOW TO CITE THIS FRAMEWORK

> *Ivezic, M. (2026). The Applied Quantum PQC Migration Framework: A Practitioner's Methodology for Enterprise Post-Quantum Cryptography Migration (Version 2.1). PostQuantum.com / Applied Quantum.*

This framework is published under CC BY 4.0. Organizations and consultants are free to use, adapt, and build on this work, including for commercial purposes, provided they credit Marin Ivezic and Applied Quantum, link to the license, and indicate any changes made. Adding proprietary copyright notices or redistribution restrictions to derivative works violates the license terms. Full attribution requirements and the framework's publication history are at PQCFramework.com/license.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

RELATIONSHIP TO OTHER GUIDANCE

This framework was first drafted in March 2023 (v0.1) and first published in full in June 2025 (v1.0). It is the first published PQC migration methodology that covers the complete migration lifecycle at operational depth in a single integrated framework. A comprehensive survey of 80+ published PQC frameworks, available at PQCFramework.com/research, found that no other single framework covers this full scope; organizations must typically combine four or five separate frameworks to assemble a coherent migration program. This framework is informed by, but not derivative of, the following standards and guidance:

- NIST FIPS 203/204/205 (ML-KEM, ML-DSA, SLH-DSA algorithm standards)
- NIST IR 8547 (Transition to Post-Quantum Cryptography Standards)
- NIST SP 1800-38 (NCCoE Migration to Post-Quantum Cryptography)
- NIST SP 800-227 (Recommendations for Key-Encapsulation Mechanisms)
- NSA CNSA 2.0 (Commercial National Security Algorithm Suite)
- PQCC Migration Roadmap (Post-Quantum Cryptography Coalition / MITRE)
- Dutch PQC Migration Handbook (AIVD/CWI/TNO, 2nd edition)
- GSMA PQ.01–PQ.03 v2.0 (Post-Quantum Telco Network Task Force)
- ETSI TR 103 619 / TR 104 016 (Migration to Post-Quantum Cryptography)
- PKI Consortium PQCMM (Post-Quantum Cryptography Maturity Model)

Where this framework takes positions that differ from conventional wisdom or other published guidance, those positions are made explicit and defended with evidence. Key differentiators include: the Minimum Viable CBOM model, the risk-driven discovery scoping approach, the emphasis on vendor governance as the primary external constraint on migration timelines, and the integrated operational security architecture (SOC and GRC implementation) that no other published PQC migration framework provides.

TOOL AND VENDOR REFERENCES

This framework references specific tools, products, and vendors as illustrative examples to help practitioners understand the categories of capability available. A mention does not constitute an endorsement, recommendation, or assessment of fitness for any particular environment. The PQC tooling market is evolving rapidly; products mentioned may have changed in capability, licensing, or availability since publication. Organizations

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

should conduct their own evaluation based on their specific requirements, regulatory environment, and procurement constraints.

ACCOMPANYING RESOURCES

Every aspect of this framework, from executive business cases and cryptographic inventory methodologies to CBOM architecture, hybrid deployment patterns, and vendor governance has been analyzed in detail on PostQuantum.com over the past 10+ years through practitioner-grounded articles, deep dives, and sector-specific analyses.

The best starting point is the curated Getting Started with Quantum Security and PQC Migration page: https://postquantum.com/starting-pqc-quantum-security/, but readers should also use PostQuantum.com's Search function to surface additional relevant posts that may not be included in that curated list.

Key framework resources, templates, and supporting materials are also published on PQCFramework.com, a dedicated companion site for this framework, drawing on relevant content from PostQuantum.com.

This framework is the execution methodology. For the complete treatment (the reasoning behind each phase, extended case examples, and guidance for leading the program from the first board conversation through closure), see the companion book, *Quantum Ready* (QuantumReady.com).

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

TABLE OF CONTENTS

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

HOW TO USE THIS FRAMEWORK

This is an executable methodology for enterprise PQC migration, not a whitepaper, not a checklist, not a vendor pitch. It is structured as an 8-phase lifecycle (Phase 0 through Phase 7) with cross-cutting concerns woven throughout. Each phase defines:

- **Prerequisites** — what must be in place before you start
- **Activities** — what you do, with enough specificity to assign to a team next Monday
- **Decision Points** — where choices branch based on your context
- **Outputs** — what you produce, with quality criteria
- **Interdependencies** — what feeds into and out of this phase
- **Common Failures** — what goes wrong and how to avoid it
- **Maturity Indicators** — how to assess whether you've done this phase well enough to proceed

Organizations should not expect to execute phases sequentially in a clean waterfall. Real programs overlap: you will be running Phase 1 discovery on some systems while executing Phase 5 pilots on others. The phase structure provides a logical dependency order; the calendar comes from the Phase 4 roadmap.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Role-Based Reading Paths

This framework runs long because the program it describes is large. No single reader needs all of it at once. The paths below identify what each role should read first; the rest is reference material for when the corresponding work begins.

CISO and executive sponsor. Phase 0, the Regulatory Timeline Context, the Maturity Levels, the board-level KPI pack in Metrics, KPIs & Reporting, the GRC Implementation foundation, and Appendix F for the objections you will hear in the funding conversation.

Program manager (QRPM). Phase 4 in full, the governance structure in Phase 0, the Year-1 quarter-by-quarter plan, the Interdependencies subsection of every phase, the Accelerated Execution Profile (Activity 4.7), and the Metrics section.

Security architect and cryptographic engineer. Phases 1, 2, 5, and 6, Appendix A for algorithm parameters, Appendix H for protocol-family coverage, and the Crypto-Agility foundation.

Risk and compliance officer. Phase 3, the GRC Implementation foundation, the Regulatory & Standards Alignment Map, the Migration Verification & Program Closure section, and Appendix G for mapping this framework onto frameworks the organization has already adopted.

Procurement and vendor management. Phase 7 in full, with particular attention to the contract language in Activity 7.3, the readiness questionnaire in Activity 7.2, and the cloud and SaaS classification in Activity 7.7.

Right-Sizing the Framework

The framework was validated on programs exceeding 120,000 tasks, but the methodology scales down. What changes by organization size is which activities merge and how heavy the governance machinery needs to be, not which questions get answered. Team sizing by CBOM instance count appears in the Skills & Team Structure foundation; the profiles below calibrate the rest.

Profile	Typical shape	What compresses or merges	Non-negotiable at any size
Small (under ~2,000 employees)	Hundreds of certificates; estate dominated	Phases 0 and 4 merge into a single charter-plus-plan; one combined inventory/CBOM workstream; SteerCo folds	Risk-driven scoping (1.0); a queryable CBOM; questionnaires to the top 10 vendors;

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

	by SaaS and cloud	into an existing security forum; one or two pilots	hybrid TLS on internet-facing endpoints
Mid-size (~2,000–20,000)	Thousands of certificates; mixed cloud and on-premises; some OT or embedded	Eight workstreams consolidate to four or five; Year-1 plan runs on a two-quarter rhythm; Waves 1 and 2 combine	A dedicated QRPM; Phases 1–3 at full Tier-1 depth; milestone gates; vendor governance cadence
Large (~20,000–100,000)	Tens of thousands of certificates; significant OT, legacy, and vendor estate	The framework applies as written	Everything, including the PMO discipline in Activity 4.4
Hyperscale or multi-entity group (100,000+)	Federated estates across subsidiaries, jurisdictions, and regulators	Nothing compresses; Phases 1–3 run as per-entity instances feeding a federated CBOM, with group-level Phase 0 and Phase 4	Federation governance; cross-entity dependency mapping; jurisdiction-specific algorithm tracks

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

FRAMEWORK ARCHITECTURE AT A GLANCE

Phases are logically sequential but operationally overlapping: real programs run discovery, documentation, pilots, and vendor engagement concurrently.

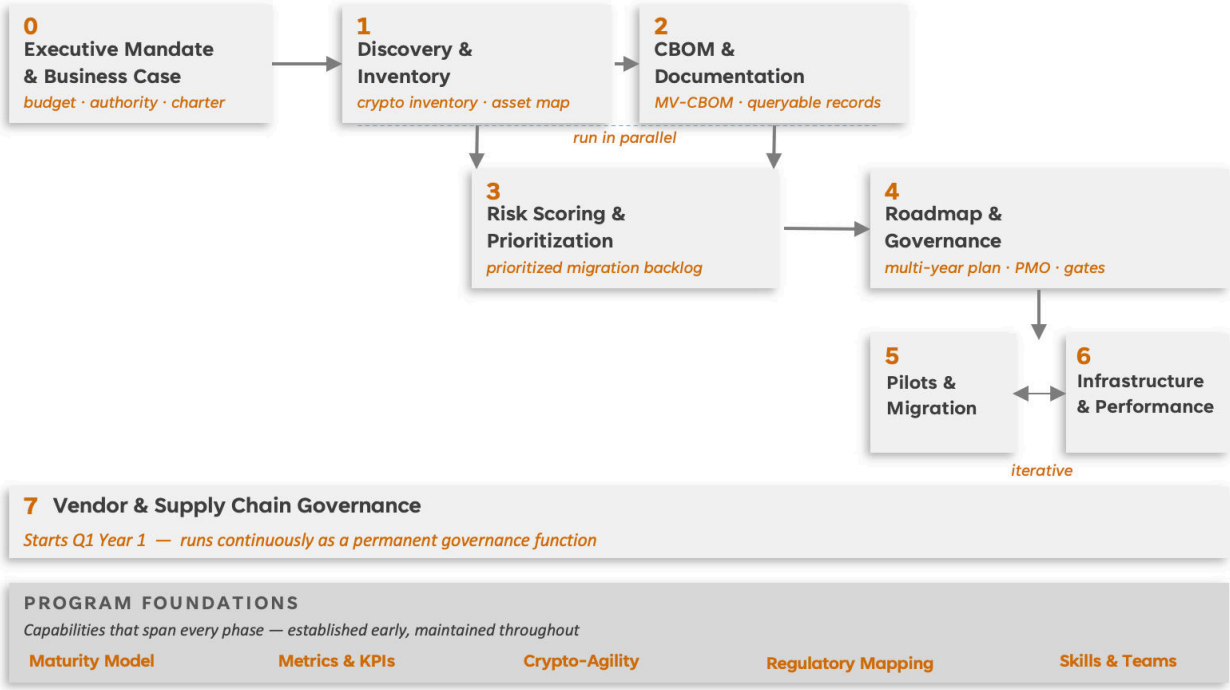

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

REGULATORY TIMELINE CONTEXT

This framework is designed to satisfy the following converging deadlines. Note: the table below includes a mix of binding requirements, recommended roadmaps, and planning milestones at varying stages of finalization. Not every date is a legally binding deadline; readers should verify the current status and applicability of each mandate for their specific jurisdiction and sector.

Jurisdiction / Body	**Key Milestone**	**Deadline**	**Status**
NIST IR 8547	Deprecate quantum-vulnerable public-key algorithms (112-bit security)	After 2030	Initial Public Draft (Nov 2024)
NIST IR 8547	Disallow quantum-vulnerable public-key algorithms	After 2035	Initial Public Draft (Nov 2024)
NSA CNSA 2.0	New NSS acquisitions CNSA 2.0 compliant	2027	Binding for NSS
NSA CNSA 2.0	Software/firmware signing uses PQC	2030	Binding for NSS
NSA CNSA 2.0	Networking equipment (VPNs, routers) exclusively uses CNSA 2.0	2030	Binding for NSS
NSA CNSA 2.0	Most NSS (web, cloud, OS, and other platforms) migrated	2033	Binding for NSS
NSA CNSA 2.0	All NSS including custom/legacy systems fully migrated	2035	Binding for NSS
Australia (ASD)	Cease traditional asymmetric crypto in Australian Government systems	2030	Government guidance
NCSC UK	Discovery and planning complete	2028	National guidance

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

NCSC UK	High-priority migration complete	2031	National guidance
NCSC UK	Full migration complete	2035	National guidance
EU Coordinated Roadmap	First steps (awareness, inventories, pilots)	End of 2026	Roadmap for Member States
EU Coordinated Roadmap	High-risk migration complete	End of 2030	Roadmap for Member States
EU Coordinated Roadmap	Medium-risk migration complete	End of 2035	Roadmap for Member States
Government of Canada (CCCS)	Departmental PQC migration plans developed	April 2026	Federal government
Government of Canada (CCCS)	High-priority systems migrated	2031	Federal government
PCI DSS v4.0 Req 12.3.3	Documentation and annual review of cryptographic cipher suites and protocols in use	Required after March 31, 2025	Binding for PCI-scoped entities

The implication: organizations starting now have 4–9 years depending on jurisdiction and system criticality. A credibly planned migration for a large enterprise requires 4–15 years of execution. There is no slack in the schedule.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PHASE 0 — EXECUTIVE MANDATE & BUSINESS CASE

PURPOSE

Nothing happens without budget, authority, and organizational commitment. Phase 0 establishes the governance foundation and secures multi-year funding. This is the single most common failure point. Programs that skip or underinvest in Phase 0 stall within 6–12 months when they hit their first resource conflict or political obstacle.

Phase 0 is also where the program's identity is established. PQC migration is frequently misunderstood as a narrow technology upgrade: a library swap or a certificate rotation project. In reality, it is the largest cryptographic overhaul most organizations will ever undertake, touching every application, every integration, every vendor relationship, and every data store that relies on public-key cryptography. The executive mandate must frame this accurately: this is a multi-year enterprise transformation program, not an IT project. Organizations that frame it as a project get project-level funding and project-level authority, which is insufficient for the scope of work ahead.

The companion book, *Quantum Ready* (QuantumReady.com), treats the executive mandate problem at chapter length, including the board conversation and the extended business case material this methodology deliberately compresses.

Parallelization note

Phase 0 is the only phase that must be substantially complete before others begin. You cannot conduct discovery without budget, tooling authority, and designated staff. However, experienced organizations often run a lightweight "Phase 0.5" concurrently with later Phase 0 activities: while the full business case and governance structure are being finalized, a small technical team can begin preliminary scoping (identifying the top 10–20 critical systems and the top 10 vendor dependencies) to generate early data points that strengthen the business case itself. This creates a virtuous cycle: early technical findings provide the evidence that secures broader executive commitment.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PREREQUISITES

Framework prerequisites (from earlier phases)

None. This is the entry point.

Organizational prerequisites

- **Awareness that quantum computing poses a material risk to the organization's cryptographic infrastructure**. This is assumed for this framework's audience. If executive awareness does not yet exist, invest in education first.
- **Access to senior leadership (CISO, CIO, or equivalent) who can sponsor a board-level initiative**. The sponsor must have the organizational standing to secure multi-year budget commitments and cross-functional authority. A mid-level security manager cannot drive this program. It requires someone who can convene business unit leaders, negotiate with procurement, and report to the board.
- **Basic understanding of the organization's regulatory environment and any sector-specific quantum readiness guidance**. At minimum, the team entering Phase 0 should know which regulatory bodies govern their sector, whether any quantum-specific mandates or timelines apply (see the Regulatory Timeline Context table in this framework), and what existing compliance frameworks (PCI DSS, DORA, CMMC, etc.) may intersect with PQC migration.
- **An organizational culture that can sustain long-horizon programs**. This is less a checklist item than a reality check. Organizations that routinely cancel or defund programs after 12–18 months will struggle with PQC migration unless Phase 0 explicitly addresses funding durability. If your organization lacks this muscle, the governance structure designed in this phase must compensate. For example, by embedding PQC migration into an existing enterprise risk management cycle that has established board-level reporting cadence.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ACTIVITIES

0.1 Frame the Business Case

Structure the executive argument around four "urgency drivers": not Q-Day predictions, which are speculative, but business pressures that are concrete and current:

1. **Regulatory and compliance deadlines.** Map the organization's specific regulatory exposure to the timeline table above. For financial services: PCI DSS v4.0 Requirement 12.3.3 already requires documentation and annual review of all cryptographic cipher suites and protocols in use, effectively mandating a cryptographic inventory. For organizations with EU operations: NIS2 and DORA create supervisory expectations. For government suppliers: CNSA 2.0 timelines apply to procurement requirements.

2. **Harvest Now, Decrypt Later (HNDL) exposure.** Identify data categories with long confidentiality requirements (10+ years): trade secrets, M&A plans, health records, national security information, long-lived financial instruments, attorney-client privileged communications. Every day these data traverse quantum-vulnerable encryption channels, adversaries can capture and store ciphertext for future decryption. This risk does not wait for a quantum computer; it is active now.

3. **Trust Now, Forge Later (TNFL) exposure.** Identify long-lived digital signatures and trust anchors: PKI root certificates (often 20+ year validity), code-signing keys, firmware signing keys, legal/regulatory document signatures, OT safety certificates. A future quantum computer could forge signatures on these artifacts, undermining trust retroactively.

4. **Client, investor, and insurer expectations.** Board directors face fiduciary duty questions. Institutional investors are incorporating quantum readiness into due diligence. Cyber insurers are beginning to include quantum preparedness in underwriting criteria. Government and enterprise procurement increasingly requires demonstrated quantum readiness.

Decision Point: If your organization holds data with >10-year confidentiality requirements AND operates in a regulated industry, the HNDL argument alone justifies immediate action. If your primary risk is signature integrity (e.g., OT/ICS environments, PKI operators), the TNFL argument drives urgency on different systems.

0.2 Build the Budget Structure

Quantum readiness is not a single project with a single budget line. Structure funding as a program with phased investment:

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Year 1 (Foundation): Inventory and discovery tooling, initial CBOM development, 2–3 hybrid pilots, vendor engagement, policy updates, team training. Illustrative range for a large enterprise: $1.5M–$4M depending on estate complexity and tooling choices; actual costs vary widely by organization size, sector, and existing tooling maturity.

Years 2–3 (Bulk Migration): Rollout to Tier-1 systems, PKI modernization, key-lifetime reductions, dual-signing pilots, vendor upgrade coordination. This is the most capital-intensive phase.

Years 4–5 (Long Tail and Hardening): Embedded/OT systems, legacy system containment, elimination of compensating controls, crypto-agility maturation.

Key budget strategy: Align PQC migration spending to existing infrastructure refresh cycles (data center refreshes, SD-WAN upgrades, cloud migrations, PKI renewals, vendor contract renewals). This avoids "big-bang" budget spikes and makes quantum readiness an incremental cost on already-planned expenditures rather than a net-new program.

The "umbrella program" strategy: Position quantum readiness as the organizing framework for long-overdue security modernization. Many organizations need to update PKI infrastructure, improve certificate lifecycle management, modernize HSMs, improve asset inventories, and strengthen vendor governance regardless of quantum risk. Quantum readiness provides the urgency and the budget justification to do all of this under one program umbrella, with PQC migration as the binding thread.

Align with innovation, R&D, and digitization budgets. Quantum readiness has significant overlap with other strategic investment categories. If the organization has an innovation or R&D budget for quantum technology exploration (quantum computing applications, quantum sensing, quantum networking), the security workstream should be funded alongside (and ideally integrated with) that quantum exploration program. The upskilling requirements overlap substantially: teams learning about quantum computing for business applications also need to understand quantum risk. Similarly, digital transformation and IT modernization programs (cloud migration, zero trust architecture, DevSecOps maturity) share infrastructure and tooling with PQC migration. Position PQC as a mandatory security track within these broader programs rather than competing for separate budget allocation. This approach also ensures that digitization initiatives do not inadvertently create new quantum-vulnerable attack surface even as they modernize the estate.

0.2b Build the Business Case — Additional Benefit Arguments

Beyond the four urgency drivers, several additional arguments strengthen the business case:

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Regulatory trust and approval acceleration. In regulated industries (financial services, healthcare, telecoms, energy), organizations frequently require regulatory approval before launching new products, services, or technology platforms. Regulators are more likely to approve new initiatives from organizations that demonstrate strong security fundamentals. An organization that can present a mature quantum readiness posture, with a documented CBOM, risk assessment, and migration roadmap, signals to regulators that it manages technology risk proactively. This can strengthen the organization's credibility in regulatory engagement and may reduce friction during supervisory review of new products and services, depending on sector, jurisdiction, and the regulator involved. In financial services, for example, demonstrating quantum readiness during a new product approval process differentiates the organization from competitors who have not addressed the risk, potentially making the difference between approval and extended review cycles.

Immediate security value from cryptographic inventory. One of the most compelling near-term benefits of PQC migration is that it produces tangible security improvements from day one, before any PQC algorithm is deployed. A quantum-vulnerability-focused cryptographic inventory inevitably also surfaces classically vulnerable cryptography: deprecated TLS versions (TLS 1.0/1.1 still in production), weak key sizes (RSA-1024, DH-1024), insecure cipher suites (RC4, 3DES, export ciphers), expired or misconfigured certificates, hardcoded keys in application code, unpatched cryptographic library versions with known CVEs, and misconfigured protocols allowing downgrade attacks. These are real, present-day vulnerabilities that can be remediated immediately. Cryptographic inventories routinely uncover classically vulnerable configurations that had not been centrally tracked or prioritized for remediation: deprecated protocols, weak keys, expired certificates, and misconfigured settings that represent real, present-day risk. This makes the Phase 1 investment self-funding from a security perspective: it pays for itself in reduced classical risk before the quantum migration even begins. Frame this in the budget request: "The inventory alone will identify and enable remediation of current cryptographic vulnerabilities, reducing our classical attack surface while simultaneously preparing for quantum risk."

Competitive differentiation and market access. By the late 2020s, demonstrating quantum resilience will become a market signal of forward-thinking security. Government and enterprise procurement requirements now include quantum readiness criteria. Organizations that can respond to "Are you quantum-ready?" with documented evidence gain competitive advantage in sales cycles, particularly in financial services, healthcare, defense, and critical infrastructure sectors.

Security talent attraction and retention. PQC and crypto-agility expertise is scarce, and demand is rising on regulatory schedules. Organizations running a serious quantum

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

readiness program offer their security engineers frontier work that competitors cannot, which functions as a recruiting and retention asset in a market where skilled cryptographic engineers choose their employers. The program's training investment compounds the effect: staff trained on PQC become more valuable, and the organization keeps them by being the place that gives them the program to run.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

0.2c Develop Migration Cost Estimates

PQC migration cost estimation is inherently imprecise -- no two organizations have the same cryptographic estate, vendor dependency profile, or regulatory environment. But "we cannot precisely estimate the cost" is not a valid reason to avoid building a budget structure. Organizations that wait for precise cost data before securing funding will wait indefinitely.

Cost Driver Taxonomy

PQC migration costs fall into eight categories. For initial budget construction, estimate each category separately:

- **Discovery and inventory tooling.** Automated cryptographic discovery platforms (SandboxAQ, IBM Quantum Safe, Keyfactor, ISARA Advance, CryptoNext, or equivalent), deployment and integration effort, ongoing license or subscription costs.
- **Cryptographic engineering labor.** Internal and external cryptographic architects, security engineers with PQC expertise, application security leads performing code and configuration changes. This is typically the largest cost category. PQC-specialized talent is scarce and commands premium rates.
- **Vendor PQC licensing and upgrades.** Vendor-controlled systems may require paid upgrades, new license tiers, or hardware replacement to gain PQC support. HSM firmware upgrades alone can cost $50K-$500K+ depending on the number of modules.
- **HSM and hardware refresh.** Older HSM models may not support PQC key types and require replacement. HSM procurement lead times are 6-12 months. Budget for both hardware cost and the significant operational effort of HSM commissioning, key ceremony, and migration.
- **PKI modernization.** Certificate authority infrastructure upgrades, automated certificate lifecycle management platforms, root CA ceremonies, intermediate CA provisioning. This investment is required regardless of PQC (shorter certificate lifetimes demand it independently).
- **Performance and capacity uplift.** PQC operations are computationally more intensive and generate larger payloads. Some systems will require CPU, memory, bandwidth, or storage capacity increases.
- **Testing environments.** Production-mirror environments for pilot validation, performance benchmarking, and compatibility testing.
- **Program management overhead.** Quantum Readiness Program Manager, PMO tools, SteerCo and working group facilitation, training programs, vendor management, regulatory compliance tracking, board reporting.

The Infrastructure Modernization Umbrella

The most effective budget strategy positions PQC migration within a broader infrastructure modernization program. Multiple investments required for PQC readiness are independently justified. PKI automation is required for 47-day certificate lifetimes regardless of PQC. FIPS 140-3 module upgrades are required for the FIPS 140-2 sunset.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

HSM refresh addresses end-of-life hardware. Library and platform upgrades remediate known CVEs.

By combining these into a single "cryptographic infrastructure modernization" program, the PQC-specific incremental cost is a fraction of the total, and the program delivers tangible security improvements from day one. This framing transforms PQC from a speculative risk investment into a concrete modernization program with immediate operational benefits and future quantum resilience as an additional outcome.

Reference Program Economics

Cost categories tell you what to estimate; the shape of real programs tells you what to expect. The anchors below are drawn from a published full-program description for a large global telecommunications operator (PostQuantum.com) and generalize in shape, if not in absolute scale, to most large enterprises:

Distribution across the lifecycle. Discovery and inventory runs in the low single-digit millions for a complex estate ($2–5M for a large operator, including tooling and a dedicated team for roughly a year). Assessment and planning is small (on the order of 1–2% of total program cost) but determines whether the remaining 98% is spent in the right order. Implementation dominates: hardware refresh, vendor upgrades, and engineering labor account for the large majority of total spend. Steady-state operations after migration settle at a few million per year.

Order of magnitude. A large, complex global enterprise should expect total program cost in the hundreds of millions over a decade ($300–500M was the modeled range for a major global telco) against the U.S. federal civilian comparator of $7.1B for 2025–2035 (OMB). Mid-size organizations scale down with estate complexity rather than headcount; the per-CBOM-instance staffing heuristic in Skills & Team Structure is the better scaling basis.

Peak staffing shape. Expect a peak of dozens of dedicated FTEs during the bulk implementation years (the telco model peaked around 50), tapering into the permanent capability described in Program Foundations.

Incremental versus reallocated spend. A substantial share of program cost is planned refresh redirected to quantum-safe equipment rather than net-new money: hardware replaced at end-of-life with PQC-capable models was budget that would have been spent anyway. Present gross program cost and net incremental cost separately; the gap between them is the strongest single number in the budget conversation, and it is the quantitative expression of the infrastructure modernization umbrella above.

These figures are credibility anchors, not estimates for your organization. Build the estimate from the cost categories above; use the anchors to defend its plausibility.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

0.3 Establish Governance Structure

Role	Responsibility	Reporting
Executive Sponsor (CISO or CIO)	Visible owner; clears roadblocks; briefs board quarterly	Board / Risk Committee
Steering Committee (SteerCo)	Cross-functional: Security, Enterprise Architecture, AppDev, Infrastructure/NetSec, PKI/Identity, Compliance/Legal, Procurement, Business Unit reps	Monthly to Sponsor
Quantum Readiness Program Manager (QRPM)	Day-to-day leader; runs plan, risk log, KPIs; coordinates workstreams	Weekly to SteerCo lead
Workstream Leads (one per domain)	Execute phase activities within their domain	Weekly to QRPM

Workstream structure (8 streams):

1. Inventory & Discovery (Crypto-BOM ownership)
2. Network & TLS/VPN (hybrid rollouts)
3. PKI & Code Signing (roots, issuers, toolchains)
4. Applications & Platforms (libraries, service mesh, cloud)
5. Embedded/IoT/OT (gateways, compensating controls)
6. Policy/Compliance/Procurement (standards, clauses)
7. Vendor Orchestration (roadmaps, SLAs)
8. Education & Change Management (training, comms)

Decision cadence:

- Weekly PMO: track milestones, blockers, vendor responses
- Monthly SteerCo: approve roadmap changes, funding asks, target dates, risk acceptance
- Quarterly Board/Risk Committee: KPIs, exceptions, budget status

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Program Leadership: Who Should Own It

The most common governance question is whether the PQC migration program should report to the CISO, the CIO, or a joint structure. The answer depends on organizational context, but the evidence from programs that have succeeded and those that have stalled points to a pattern.

The CISO should lead the program in most organizations. PQC migration is a security risk response, not a technology refresh. The CISO owns the threat model (HNDL, TNFL), the risk register entry, the regulatory compliance obligations, and the relationship with the board's risk committee. Programs that report to the CIO tend to be framed as infrastructure modernization, which is accurate at the execution layer but misses the risk governance that drives prioritization. A CIO-led program will optimize for operational efficiency; a CISO-led program will optimize for risk reduction. Both matter, but risk reduction determines sequencing.

The CIO's role is essential but different. The CIO owns the technology estate being migrated, the vendor relationships that constrain timelines, the infrastructure budget, and the enterprise architecture function that determines whether crypto-agility gets built into new systems. The CIO should co-sponsor the program and chair the technical workstreams, but the CISO should own the program charter, the risk appetite, and the board reporting relationship.

Joint structures work when the organization already has a mature operating model for shared CISO-CIO programs (such as a joint security-and-technology risk committee). They fail when "joint" means "neither has clear accountability." If choosing a joint model, designate one as the accountable sponsor for board reporting and escalation decisions.

In organizations where the CISO reports to the CIO, the program should still be chartered as a risk-driven initiative with a direct reporting line to the board's risk committee or audit committee. This reporting line matters: PQC programs buried inside IT portfolios compete for priority against ERP upgrades, cloud migrations, and other infrastructure projects that have more immediate operational urgency. A direct risk committee reporting line gives the program the organizational gravity to survive leadership changes and budget cycles.

Board Oversight

Board oversight should operate through the existing risk committee or audit committee. Creating a separate "quantum committee" marginalizes the program rather than integrating it. PQC migration is a risk management activity and should compete for attention through the same governance channel as cyber risk, operational risk, and compliance risk.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Minimum board engagement model:

Initial briefing (60–90 minutes, conducted once during Phase 0). Cover the quantum threat in business terms: HNDL and TNFL exposure mapped to the organization's data and trust infrastructure, regulatory deadlines that create compliance risk, fiduciary duty implications of inaction when standards and solutions exist, estimated program scope and cost, and the proposed risk appetite statement. The goal: an informed board that can approve a risk appetite statement and a multi-year budget commitment.

Quarterly KPI review (15–20 minutes within existing risk committee agenda). Five KPIs, reported as a single-page dashboard. These are the board-level KPIs defined in the Metrics, KPIs & Reporting section: Coverage, Trust, Inventory, Vendors, and Agility. The board should see the number, the trend, the target, and an explanation for any variance exceeding 5 percentage points.

Annual risk appetite review. The board reviews and re-approves the quantum risk appetite statement (described below), incorporating changes in the threat landscape, regulatory environment, and migration progress.

Escalation triggers (immediate board notification outside the quarterly cycle). A material change to the CRQC timeline estimate from a credible source. The migration program falling more than 6 months behind its regulatory compliance buffer. A confirmed vulnerability in a deployed PQC algorithm. A Tier 1 vendor abandoning PQC support without an alternative.

Risk Appetite Statement

A quantum risk appetite statement translates the board's intent into decision criteria for the program. Without it, the program manager has no anchor when facing trade-offs between migrating a complex legacy system (expensive, disruptive) and deferring it (cheap, smooth, but leaves a gap).

The statement operates at two levels:

Strategic level. "The organization will complete migration of all systems protecting data with confidentiality requirements exceeding 10 years before the earliest credible CRQC estimates, and will achieve compliance with all applicable PQC regulatory deadlines with a minimum 12-month buffer."

Operational tolerances:

HNDL exposure. "No more than 20% of data classified as having secrecy requirements exceeding 10 years will remain protected by quantum-vulnerable algorithms by end of 2027. The target is 0% by end of 2029."

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

TNFL exposure. “All production software and firmware signing will use NIST-approved quantum-resistant signatures (SP 800-208 hash-based or ML-DSA, dual-signed with classical) by end of 2027. All certificate authority signing keys will be PQC-capable, with hybrid or PQC certificate issuance available, by end of 2029.”

Regulatory compliance. “The organization will achieve compliance with all PQC-related regulatory requirements at least 12 months before the applicable deadline. Zero tolerance for deadline non-compliance.”

Crypto-agility. “All newly deployed systems from Q3 2026 onward must implement crypto-agile architecture as a mandatory design requirement. No exceptions without Steering Committee approval and a documented remediation plan.”

These thresholds are illustrative. The right values depend on the organization’s sector, regulatory environment, data profile, and risk culture. The requirement is that they exist, are specific enough to drive decisions, and are reviewed annually.

Three Lines of Defense

PQC migration maps naturally to the three-lines-of-defense model that most regulated organizations already use:

First line (business and technology teams) owns the cryptographic estate and executes the migration. Application teams migrate their code. Infrastructure teams upgrade PKI, HSMs, and network devices. Procurement updates contract clauses. Each workstream lead is a first-line function. The first line does the work.

Second line (CISO / risk / compliance) owns the governance framework, sets the risk appetite, monitors KPIs, manages the regulatory intelligence function, and reports to the board. The QRPM sits here. The second line does not execute migration tasks; it governs the program, measures progress, and escalates when first-line execution falls behind plan. The second line also owns the operational risk instruments: cascading KRIs, the regulatory horizon tracker, and third-party quantum readiness assessment.

Third line (internal audit) provides independent assurance that the program is governed as chartered, the cryptographic inventory is accurate, the KPIs are measured honestly, and the migration is actually being executed (not just reported as executed). Recommended audit timing: initial audit within 6 months of program launch (governance and inventory validation), mid-program audit at 18–24 months (execution against plan), annual thereafter until migration is substantially complete. Internal audit should also observe tabletop exercises to assess incident preparedness.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Cascading Key Risk Indicators

Risk indicators for PQC migration should cascade across three organizational levels. Each level has a different audience, reporting frequency, and level of detail. The full KRI framework with illustrative thresholds is specified in the GRC Implementation section of the Program Foundations.

Operational Security Integration

PQC migration creates operational responsibilities for two functions that most frameworks have not addressed: the Security Operations Center (SOC) and the Governance, Risk, and Compliance (GRC) function. These responsibilities extend beyond the migration program; they become permanent organizational capabilities.

The SOC must develop detection capabilities for five quantum-related threat scenarios: hybrid downgrade detection, cryptographic drift monitoring, certificate lifecycle anomalies, TNFL and signature integrity monitoring, and enhanced HNDL-indicator detection. The SOC also houses the Cyber Threat Intelligence (CTI) function responsible for tracking CRQC developments, monitoring PQC implementation vulnerabilities, and producing the quarterly Quantum Threat Landscape Assessment that feeds into the Material Developments KRI at board level.

The GRC function owns the risk appetite statement, the KRI framework, the regulatory intelligence function (quarterly Regulatory Horizon Report tracking pending requirements across all jurisdictions), third-party quantum readiness assessment (described in Phase 7), the evidence dossier for audit (described in Metrics, KPIs & Reporting), and insurance preparation.

The critical handoff between SOC and GRC is the cryptographic inventory. The migration program builds it. GRC governs it and uses it for compliance reporting. The SOC needs it in machine-readable, SIEM-integrated form to make detection rules function. An inventory that lives in a quarterly-updated spreadsheet serves audit but not detection. Designing the inventory as a shared operational data asset should be a Phase 1 architecture decision.

The full operational specifications for SOC detection capabilities, CTI requirements, incident response playbooks, tabletop exercises, and implementation roadmap are documented in the SOC Implementation section of the Program Foundations. The corresponding GRC instruments (the KRI framework, regulatory intelligence process, vendor assessment mechanics, audit procedures, insurance preparation, and the GRC-SOC handoff) are documented in the GRC Implementation section.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

0.4 Draft the Program Charter

A one-page charter document covering:

- **Purpose:** Quantum-safe migration and crypto-agility
- **Scope:** TLS/VPN, PKI/code-signing, applications/platforms, embedded/OT, policy/procurement, vendors, training
- **Success criteria:** Stop new HNDL exposure; protect long-lived trust anchors; meet 2030/2035 regulatory timelines; embed crypto-agility as organizational capability
- **Cadence:** Weekly PMO; monthly SteerCo; quarterly board
- **Escalation path:** Sponsor decides on funding/dates; risk acceptance decisions are documented and auditable

0.5 Conduct Initial Scoping Assessment

Before launching full discovery, perform a rapid (2–4 week) scoping assessment to establish program boundaries:

1. Identify the top 20 revenue-generating or mission-critical systems
2. For each, determine: primary cryptographic protocols in use, data sensitivity classification, number of dependent systems, vendor ownership vs. internal control
3. Estimate the approximate size of the cryptographic estate (number of TLS endpoints, certificates, VPN tunnels, HSM-protected keys, code-signing pipelines)
4. Identify the 5–10 vendors whose PQC readiness will most constrain the migration timeline

This scoping assessment becomes the input for Phase 1 prioritization and helps calibrate Year 1 budget and staffing.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

OUTPUTS

Output	Quality Criteria
Approved program charter	Signed by executive sponsor; reviewed by SteerCo
Multi-year budget commitment	Minimum 3-year funding approved; aligned to refresh cycles
Governance structure	SteerCo membership confirmed; QRPM appointed; meeting cadence established
Initial scoping assessment	Top 20 systems identified; estate size estimated; critical vendor dependencies mapped
Board briefing deck	Delivered at board/risk committee level; documented in minutes

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

INTERDEPENDENCIES

Backward dependencies

See Prerequisites above. Phase 0 is the framework entry point with no dependency on prior phases.

Feeds into

- **Phase 1** — The scoping assessment (Activity 0.5) determines discovery priorities and initial system targeting. The approved budget enables tool procurement and team staffing. The governance structure provides the authority needed to gain access to systems, repositories, and network infrastructure.
- **Phase 4** — The budget structure and multi-year commitment set the financial constraints within which the roadmap must be built. The governance model (SteerCo, RACI, escalation paths) becomes the operating framework for roadmap governance.
- **Phase 7** — The critical vendor list from the scoping assessment enables vendor engagement to begin in Q1 Year 1, before CBOM or risk scoring is complete. Procurement and legal team engagement initiated in Phase 0 provides the contracting capability needed for Phase 7.

Runs in parallel with

Late-stage Phase 0 activities (governance finalization, detailed budget approval) can overlap with early Phase 1 discovery on the highest-priority systems, as described in the Purpose section. Early technical findings from this preliminary discovery strengthen the business case, creating a virtuous cycle between Phase 0 justification and Phase 1 evidence.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

COMMON FAILURES

- **"Innovation project" framing.** Treating PQC migration as an R&D initiative or skunkworks project instead of a funded, governed program. This guarantees it will be deprioritized when competing for resources with operational imperatives.
- **Single-year budget.** Securing one year of funding and hoping to "show results" to justify the next year. PQC migration is a 4–15 year program. Without multi-year commitment, teams cannot plan realistically or retain specialized talent.
- **Missing business unit representation on SteerCo.** The migration will touch every application, every integration, every vendor relationship. Without business unit buy-in at the governance level, workstream leads face constant political resistance.
- **Delegating to vendors.** Assuming "our vendors will sort this out" and therefore not needing an internal program. This is the single most dangerous misconception. Vendors will update their products on their own timelines, optimizing for their own priorities. Without an internal program driving requirements, tracking commitments, and testing deployments, the organization has no control over its migration timeline.
- **Premature lockdown.** Reacting to quantum headlines or a regulator's inquiry by abruptly restricting data access or disabling systems to demonstrate seriousness. HNDL risk is not reduced by making data harder for your own organization to use: already-harvested traffic is already gone, and knee-jerk restrictions burn the operational goodwill the program will need for a decade. Channel the urgency into the structured program instead.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

MATURITY INDICATORS

Level	Indicator
Level 0 — Unaware	No executive awareness of quantum risk; no budget discussion
Level 1 — Aware	Quantum risk acknowledged; no formal program
Level 2 — Initiated	Charter approved; QRPM appointed; Year 1 budget secured
Level 3 — Established	Multi-year budget committed; SteerCo operational; scoping assessment complete
Level 4 — Optimized	Program integrated into enterprise risk register; quantum risk reported to board quarterly alongside other strategic risks

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PHASE 1 — DISCOVERY & INVENTORY

PURPOSE

Build a comprehensive, continuously updated inventory of all cryptographic usage across IT, OT, cloud, IoT, and third-party systems. This is the foundational activity upon which every subsequent phase depends. You cannot migrate what you cannot see, and you cannot prioritize what you have not inventoried.

Major organizations should plan for this phase to take 12–24 months of dedicated team effort. The situation is further complicated by tool vendors who may imply that their products provide a complete solution. No tool provides 100% discovery on its own. A comprehensive approach combining automated tools, manual audits, and continuous monitoring is necessary.

Discovery is the phase where most organizations are currently stuck, not because the task is conceptually difficult, but because it exposes how little most enterprises actually know about their own cryptographic estate. The asset discovery problem that underpins cryptographic inventory is itself a longstanding IT governance gap: CMDBs are incomplete, shadow IT proliferates, cloud resources spin up without central oversight, and OT environments often have no digital inventory at all. Phase 1 must confront this honestly. Organizations that treat cryptographic discovery as an isolated exercise, separate from the broader challenge of knowing what they actually own, will produce inventories that are incomplete from day one.

Parallelization note

Phase 1 runs in parallel with Phase 2 from an early stage: the CBOM data structure should be defined before large-scale discovery begins, so that inventory data is collected in a format that populates the CBOM directly rather than requiring later reformatting. Phase 1 also begins to generate inputs for Phase 7 (vendor dependency data) and Phase 3 (the first risk-relevant metadata). In mature programs, Phase 1 never fully ends. It transitions from a project-mode "initial discovery" into a permanent operational capability (continuous discovery) that feeds the CBOM and risk scoring processes indefinitely. Organizations should plan staffing accordingly: the discovery team is not a temporary project team, it is a permanent function.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PREREQUISITES

Framework prerequisites (from earlier phases)

- **Phase 0 outputs: approved charter, budget, governance structure, and initial scoping assessment**. The scoping assessment from Phase 0 is particularly important: it identifies the top 20–50 systems by business criticality that form the starting scope for risk-driven discovery (Activity 1.0). Without this scoping input, discovery teams default to either "boil the ocean" (try to scan everything simultaneously) or "path of least resistance" (scan whatever is easiest), neither of which produces actionable results on a useful timeline.
- **Designated workstream lead for Inventory & Discovery**. This should be someone with cross-functional access and technical credibility, typically a senior security architect or infrastructure lead. They will need to negotiate access to systems, repositories, and network segments across business units that may not initially see PQC migration as their problem.

Organizational prerequisites

- **Access to network monitoring infrastructure, code repositories, configuration management databases (CMDBs), and certificate management systems**. In practice, gaining this access is often the first political test of the Phase 0 mandate. If the discovery team cannot get read access to network taps, source code repositories, cloud account configurations, and certificate inventories, the executive sponsor needs to intervene early. Discovery that is blocked by access politics produces coverage gaps that compound through every subsequent phase.
- **At least a preliminary asset inventory or architecture documentation**. Perfect asset data is not required. If it were, most organizations could never start. But the team needs some baseline understanding of the IT and OT estate: major application platforms, network segments, data centers, cloud accounts, OT zones. If this baseline does not exist, Phase 1 must begin with an asset discovery workstream running in parallel with cryptographic discovery, and the program plan should account for the additional time this requires.
- **Budget for cryptographic discovery tooling**. Automated discovery tools (network scanners, code analyzers, certificate crawlers) are not optional: manual-only discovery is unreliable and unscalable. Tool procurement and deployment lead times should be factored into the Phase 1 timeline. Organizations that defer tool selection until after Phase 1 "starts" lose 2–3 months to procurement and deployment before any meaningful discovery occurs.
- **Cooperation agreements with OT/ICS teams (if applicable)**. OT environments require different discovery approaches, different tools, and different risk tolerances. Scanning an OT network with IT-grade tools can cause operational disruptions. The OT team must be engaged before any discovery activities touch industrial environments, ideally during Phase 0 scoping, but no later than the start of Phase 1.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ACTIVITIES

1.0 Risk-Driven Scoping — Decide What to Inventory First

Before launching broad discovery, apply an 80/20 prioritization to determine where to focus initial effort. Attempting to discover everything simultaneously is the most common cause of inventory paralysis. The architecture-first approach (aligned with the Minimum Viable CBOM model in Phase 2) concentrates initial discovery where risk concentrates, delivering an actionable inventory quickly, then expanding systematically to comprehensive coverage.

Step 1 — Identify Tier-1 systems using existing organizational knowledge. You do not need a cryptographic inventory to know which systems matter most. Use existing sources (revenue data, business impact assessments, regulatory scope documents, incident history, architecture diagrams) to identify the top 20–50 systems by business criticality.

Step 2 — Classify by exposure type. For those Tier-1 systems, rapidly classify:

- Internet-exposed flows where adversaries can easily harvest encrypted traffic (external TLS, partner VPNs, email gateways): these face immediate HNDL risk
- Long-lived secrecy data where the confidentiality requirement exceeds 10 years (health records, trade secrets, financial instruments, legal records): these have the highest HNDL value
- Long-lived trust anchors where signature validity extends 5+ years (PKI roots, code-signing keys, firmware signing): these face TNFL risk

Step 3 — Assign discovery priority tiers.

- **Discovery Priority A:** Internet-exposed + handles long-lived secrecy data or trust anchors → Discover these first (target: 30–60 days)
- **Discovery Priority B:** Internal Tier-1 systems with significant data sensitivity → Discover next (target: 60–120 days)
- **Discovery Priority C:** Everything else → Discover on an ongoing basis (target: 6–12 months)

This approach delivers 70–80% of risk coverage with 20–30% of total discovery effort. You will identify the highest-priority migration targets quickly while comprehensive discovery continues in the background. The initial discovery output (Priority A systems) is sufficient to begin Phase 2 CBOM population and Phase 3 risk scoring on the most critical systems, while the broader inventory catches up.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

1.1 Establish Three Parallel Inventory Tracks

Quantum readiness requires integrated discovery across three domains, each with distinct methodologies but thoroughly interdependent results:

Track A — Cryptographic Inventory: Identifies and documents all uses of cryptography: algorithms, key sizes, protocols, libraries, certificate chains, and key lifetimes. This is the primary PQC-specific deliverable.

Track B — Sensitive Data Discovery and Classification: Identifies, catalogs, and classifies all sensitive data by confidentiality requirements and retention periods. This determines which cryptographic protections have the highest HNDL urgency.

Track C — Systems and Assets Inventory: Catalogs all hardware and software assets, their criticality classifications, vendor ownership, and lifecycle status. This determines migration feasibility and sequencing.

These three tracks must be coordinated through a single governance structure, even if executed by different teams using different tools. The intersection of all three ("this system (Track C) protects this sensitive data (Track B) using this vulnerable algorithm (Track A)") is what enables defensible prioritization in Phase 3.

Assigning confidentiality horizons (Track B method). Standard classification schemes grade sensitivity, not time; quantum risk scoring needs both. For each data category, assign a confidentiality horizon: the number of years the data must remain confidential, derived from, in order of precedence: (1) statutory and regulatory retention or secrecy requirements, which usually set the longest anchor (health records 10–30 years depending on jurisdiction, audit and tax records 7–10, classified material per national rules); (2) contractual confidentiality obligations, including NDA survival clauses and customer data-protection commitments; (3) intellectual property lifetime; trade secrets remain sensitive for as long as they confer advantage, frequently 15 or more years, while patent-backed material loses secrecy value at publication; and (4) business value decay for categories the first three do not govern (M&A material is typically sensitive for 2–5 years, strategic plans 3–7, pricing data 1–3). Where multiple sources apply, the longest horizon governs. Record the horizon on each data store in the classification register; Phase 3 Dimension 1 consumes it directly. Resist classifying everything as long-lived: inflated horizons destroy the prioritization signal the scoring model depends on.

1.2 Deploy Cryptographic Discovery — Layered Approach

No single discovery technique covers all cryptographic usage. Deploy a combination:

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Layer 1 — Network Traffic Analysis (Passive) Deploy passive monitoring on network taps or SPAN ports to capture TLS/SSH/IPsec handshakes and extract negotiated cipher suites, certificate chains, key sizes, and protocol versions. Configuration shows intent; this reveals what is actually negotiated in production.

Coverage: External-facing TLS, internal east-west TLS/mTLS, VPN tunnels, SSH sessions. Limitation: Cannot see encrypted payloads, application-layer cryptography, or data at rest.

Layer 2 — Static Code Analysis Scan source code repositories for cryptographic API calls, hardcoded algorithms, key generation patterns, and library imports. Target both first-party code and third-party dependencies.

Coverage: Application-layer cryptography, embedded algorithm choices, library dependencies. Limitation: Cannot detect runtime behavior, dynamically loaded crypto, or third-party binaries without source.

Layer 3 — Configuration and Certificate Scanning Enumerate all TLS certificates from certificate management systems and CT logs. Scan network device configurations (load balancers, firewalls, VPN concentrators, proxies) for cipher suite configurations. Query cloud provider APIs for KMS key metadata, managed certificate configurations, and encryption-at-rest settings.

Coverage: Certificate inventory, configured (vs. negotiated) cipher suites, cloud encryption settings. Limitation: Shows configured policy, not actual runtime negotiation.

Layer 4 — Runtime and Binary Analysis For systems without source code access (vendor appliances, legacy binaries, embedded firmware), use runtime instrumentation, binary analysis, or memory dump analysis to identify cryptographic operations.

Coverage: Vendor black-box systems, legacy applications, embedded devices. Limitation: Requires specialized skills; may not be feasible for all systems.

Layer 5 — Manual Investigation Interview application owners, review architecture documentation, examine vendor security documentation, and audit HSM/KMS usage logs. This catches cryptographic usage that automated tools miss: custom protocols, proprietary encryption, embedded hardware crypto, and undocumented integrations.

Coverage: Everything automated tools miss. Limitation: Labor-intensive; dependent on institutional knowledge.

Decision Point — Tool Selection:

The cryptographic discovery tooling market is evolving rapidly. Rather than selecting based on a static vendor list, evaluate tools across the following capability categories. An

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

effective discovery program requires coverage across multiple categories, as no single tool covers all of them:

Category 1 — Dedicated Cryptographic Discovery Platforms. Purpose-built tools that combine multiple discovery methods (network monitoring, code scanning, configuration analysis) into an integrated platform with CBOM output. These platforms are designed specifically for PQC readiness and provide the most quantum-relevant output. The current vendors include established security companies (such as SandboxAQ, IBM, Keyfactor, which acquired both InfoSec Global and CipherInsights, ISARA Advance, CryptoNext Security, and Palo Alto Networks) as well as emerging specialists. This category is evolving quickly; new entrants and capability expansions appear regularly.

Category 2 — Network Traffic Analysis and Protocol Inspection. Tools that passively capture and analyze TLS/SSH/IPsec handshakes to determine negotiated cipher suites, certificate chains, and protocol versions. Some dedicated cryptographic platforms include this capability; alternatively, network security monitoring tools and TLS inspection appliances can be repurposed.

Category 3 — Static Code Analysis (SAST) with Cryptographic Detection. Tools that scan source code for cryptographic API calls, hardcoded algorithms, key generation patterns, and library imports. Some general-purpose SAST tools can be configured with custom rules for cryptographic detection; dedicated crypto-focused scanners provide more targeted results.

Category 4 — Software Composition Analysis (SCA) and SBOM Generation. Tools that enumerate third-party library dependencies and generate Software Bills of Materials. When combined with cryptographic vulnerability databases, SCA output reveals which libraries contain quantum-vulnerable cryptographic implementations and which applications depend on them.

Category 5 — Certificate and PKI Discovery. Tools that enumerate all TLS/X.509 certificates across the estate, including internal CA-issued certificates, cloud-managed certificates, and CT log monitoring for external certificates. Certificate management platforms and PKI vendors typically provide this capability.

Category 6 — Cloud Security Posture Management (CSPM). Tools that query cloud provider APIs to enumerate encryption configurations, KMS key metadata, managed certificate settings, and storage encryption status across multi-cloud environments.

Category 7 — Binary Analysis and Reverse Engineering. Specialized tools for analyzing compiled binaries and firmware to detect cryptographic operations in vendor products where source code is unavailable. Also includes runtime instrumentation and memory

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

analysis tools. This is the most specialized category and is typically needed only for Layer 4 (embedded/third-party) discovery.

Category 8 — Extracting Cryptographic Metadata from Existing Security Tools. Some organizations can extract significant cryptographic intelligence from security tools already deployed: SIEM logs (TLS handshake data), vulnerability scanners (crypto-related CVEs), endpoint protection platforms (library version data), and network monitoring solutions (protocol version data). This approach avoids deploying new tools entirely for initial discovery and can provide rapid baseline coverage.

Category 9 — AI-Assisted Analysis and Enrichment. A tool category that emerged in 2026: large language models applied to cryptographic discovery output. Used well, these tools accelerate three specific tasks: triaging static-analysis findings (separating detected cryptographic calls that are real migration work from dead code and test fixtures), enriching CBOM entries from unstructured vendor documentation, and generating first-draft test cases for migrated components. Research published in 2026 reports fine-tuned models reaching roughly 92% functional correctness on cryptographic code migration, and LLM-assisted static analysis pipelines now pair pattern detection with model-based contextual classification. Treat every AI-generated finding, enrichment, or classification as a candidate that requires verification before it enters the CBOM as fact, and treat “AI-powered” as a marketing adjective until the vendor explains which task the model performs and how its error rate was measured. The framework’s position on AI-modified cryptographic code itself is in Activity 5.7.

For a comprehensive and current analysis of specific vendors within these categories, including capability comparisons and selection guidance, see "Cryptographic Inventory Vendors and Methodologies" on PostQuantum.com (https://postquantum.com/post-quantum/cryptographic-inventory-vendors/). Vendor capabilities shift quickly; specific product features and vendor positioning may change between publication cycles.

Recommendation: Deploy tools from at least Categories 1–2 for broad automated coverage AND supplement with Category 7/8 approaches for difficult-to-reach systems AND conduct manual investigation (Layer 5) for the top 20 critical systems identified in Phase 0 scoping. No tool alone achieves completeness. Prioritize tools that produce output in CycloneDX CBOM format or that integrate with your selected CBOM tooling from Phase 2.

1.3 Map the Cryptographic Estate

For every cryptographic instance discovered, record:

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Field	Description	Example
System/Asset ID	Unique identifier, linked to CMDB	APP-0142
Cryptographic Function	What the crypto does	Key exchange, signing, encryption at rest
Algorithm	Specific algorithm in use	RSA-2048, ECDHE-P256, AES-256-GCM
Key Size	Key length in bits	2048, 256, 384
Protocol	Transport/application protocol	TLS 1.2, SSH 2.0, IPsec IKEv2
Library/Implementation	What provides the crypto	OpenSSL 3.0.12, BoringSSL, Java 17 JCA
Certificate Details	Issuer, validity, chain depth	DigiCert G2, expires 2025-11-01, chain depth 3
Key Lifetime	How long keys persist	Session (ephemeral), 1 year, 20 years (root CA)
Data Sensitivity	Classification of protected data	Confidential, Restricted, Public
Quantum Vulnerability	Vulnerable to Shor, Grover, or neither	Shor (RSA key exchange), not applicable (AES-256)
Owner	Responsible team/individual	Platform Engineering / J. Smith
Vendor Dependency	Whether migration requires vendor action	Yes — Vendor X controls firmware
Control Posture	Whether organization controls both endpoints	Full control / Partial / No control

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

1.4 Address the Asset Discovery Problem

Cryptographic inventory depends on knowing what assets exist. Most organizations underestimate their cryptographic footprint. A typical enterprise assumes it has dozens of TLS endpoints; discovery reveals hundreds. RSA is assumed to be only in web servers. It is actually in VPN concentrators, email gateways, IoT devices, backup systems, and embedded firmware.

Why asset discovery is a prerequisite, not an afterthought. You cannot perform a cryptographic scan on a system you do not know exists. The asset discovery problem is the unacknowledged gap that underpins the entire quantum readiness program. Organizations frequently discover during Phase 1 that their existing asset registers are incomplete, particularly for cloud resources, developer-provisioned services, and OT/IoT devices. Quantum readiness forces organizations to confront this problem because the consequences of missing assets are more severe (a single undiscovered, unmitigated cryptographic endpoint is a potential quantum-era breach vector).

Comprehensive asset discovery sources: go beyond the CMDB:

The CMDB is rarely complete or current. Supplement it with every available data source:

1. **Configuration Management Database (CMDB).** The starting point, but not the ending point. CMDBs are often stale, incomplete, and biased toward officially provisioned infrastructure. Use as a baseline, then validate and extend.

2. **IT Asset Management (ITAM) databases.** Organizations often maintain separate asset management systems from their CMDB, particularly for hardware lifecycle tracking, software licensing, and financial asset registers. These may capture assets the CMDB misses, especially hardware that was procured but not formally onboarded into configuration management.

3. **Procurement and purchasing records.** Procurement systems (purchase orders, invoices, receiving records) provide a historical trail of every hardware and software acquisition. Cross-referencing procurement data against CMDB and ITAM records reveals assets that were purchased but never formally registered, a surprisingly common occurrence, especially for departmental purchases, project-specific hardware, and lab/test equipment.

4. **Cloud management consoles and APIs.** Cloud estates (AWS, Azure, GCP, and other providers) must be queried directly through their APIs and management consoles. Cloud resources are frequently provisioned outside of traditional CMDB workflows, especially in organizations with decentralized cloud access. Use cloud security posture management (CSPM) tools or native cloud inventory services (AWS Config, Azure Resource Graph, GCP Cloud Asset Inventory) to enumerate all cloud resources, including those in non-production accounts that may still process real data.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

5. **Network scanning and traffic analysis.** Active network scanning (Nmap, Shodan for external exposure, asset discovery scanners) and passive traffic analysis (NetFlow, DNS query logs, DHCP logs) reveal devices actively communicating on the network that may not appear in any register. This is particularly important for discovering rogue devices, shadow IT, and IoT endpoints.

6. **Certificate Transparency (CT) logs.** For externally visible TLS certificates, CT logs provide a comprehensive record of every certificate issued for the organization's domains. This reveals subdomains and services the organization may not have formally registered.

7. **DNS zone files and records.** Internal and external DNS records reveal hostnames and services that may not appear in asset registers. Stale DNS records can also point to decommissioned but still-reachable services.

8. **Physical walkthroughs and site surveys.** For organizations with significant OT, manufacturing, or physical infrastructure, physical walkthroughs of data centers, control rooms, factory floors, building management system closets, and network distribution rooms are essential. Physical inspection frequently reveals connected devices (cameras, sensors, building management controllers, access control panels, legacy terminals, embedded systems) that do not appear in any IT asset register because they were provisioned by facilities management, building owners, or OT teams operating independently from IT. For OT-heavy organizations (utilities, manufacturing, oil and gas, transportation), physical walkthroughs are not optional: they are the only reliable way to discover the full embedded device population.

9. **Service desk and incident records.** Tickets referencing systems that do not appear in the CMDB indicate unregistered assets. Incident response records may reference systems encountered during investigations that were previously unknown.

10. **Third-party and vendor-managed systems.** Many organizations host vendor-managed systems (managed security services, outsourced applications, co-located equipment) that do not appear in internal asset registers because the vendor maintains them. These systems still process organizational data and use cryptographic protocols that are in scope for quantum readiness.

11. **Mergers and acquisitions history.** Acquired entities frequently bring entire technology estates that were never fully integrated into the acquirer's CMDB. Post-M&A environments are notorious for hidden systems, particularly legacy infrastructure from the acquired organization that was "going to be decommissioned" but never was.

Decision Point: If your organization cannot produce a reasonably complete asset inventory (>80% coverage of IT systems, >60% of OT systems), you have a foundational problem that must be addressed before or in parallel with cryptographic discovery. Consider whether the quantum readiness program should include an asset discovery workstream, or whether it should be a prerequisite funded separately. In practice, the

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

quantum readiness program often becomes the forcing function that finally drives asset inventory completeness. This is another example of quantum readiness producing immediate security value beyond quantum-specific risk.

1.5 Integrate with Existing Data Sources

Do not build the cryptographic inventory from scratch in isolation. The organization already possesses significant relevant data across multiple systems. Integrating with these sources saves effort, improves accuracy, and ensures the cryptographic inventory inherits existing business context that would otherwise need to be re-gathered:

- **CMDB:** Asset ownership, lifecycle status, criticality classification, configuration baselines, relationships between components
- **IT Asset Management (ITAM) databases:** Hardware models, firmware versions, software licenses, asset lifecycle stage (active, end-of-life, decommissioned-but-still-running)
- **Business Impact Assessments (BIAs):** Criticality classifications, recovery time objectives (RTOs), recovery point objectives (RPOs), maximum tolerable downtime. BIAs are particularly valuable for Phase 3 risk scoring because they provide pre-existing, business-validated criticality ratings that do not need to be re-derived. If a system is classified as "Critical" in the organization's BIA, that classification should flow directly into the quantum risk scoring model rather than being re-assessed independently
- **Data classification registers:** If the organization has already classified data by sensitivity (as many regulated organizations have), this feeds directly into HNDL risk scoring without re-doing data discovery from scratch
- **Certificate Management Systems:** Certificate inventory, expiry dates, issuer relationships, certificate chain structure, renewal history
- **SBOM (Software Bill of Materials):** Library versions, dependency trees, known vulnerabilities in dependencies
- **Vulnerability Management:** Known CVEs in cryptographic libraries; historical vulnerability scan data showing which systems have had crypto-related findings
- **Cloud Security Posture Management (CSPM):** Cloud encryption configurations, KMS key policies, managed certificate inventory across multi-cloud
- **Network Security Monitoring:** Traffic flows, protocol negotiations, connection metadata

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

- **Enterprise Architecture repositories:** Application portfolio management databases, technology reference architectures, integration maps showing system-to-system connections
- **Procurement systems:** Vendor product catalogs with version information, contract records showing which vendor products are in use, license entitlements
- **Service catalogs and API registries:** Published APIs and services with their protocol and authentication requirements
- **Identity and access management (IAM) systems:** Authentication method inventory (which systems use certificate-based auth, which use SAML/OIDC with crypto dependencies)
- **Backup and disaster recovery systems:** Encryption-at-rest configurations for backup infrastructure; key management for backup encryption

The more existing data sources you integrate, the faster and more accurate the cryptographic inventory becomes. Each source fills a different gap and provides different business context that enriches the CBOM and accelerates Phase 3 risk scoring.

1.6 Establish Continuous Discovery

The cryptographic inventory is a living data set, not a one-time exercise. Cryptographic posture changes constantly: new applications are deployed, libraries are updated, certificates are issued and rotated, vendors release firmware updates, cloud configurations change, and shadow IT continues to appear. An inventory that is not continuously maintained begins to decay within weeks and becomes dangerously unreliable within months.

Continuous discovery is also a prerequisite for crypto-agility (the end-state goal described in the Cross-Cutting section). An organization that cannot maintain current visibility into its cryptographic posture cannot claim to be crypto-agile, because it cannot verify that algorithm changes have been applied or detect when systems drift from policy.

1.6.1 CI/CD Pipeline Integration

Integrate cryptographic scanning into the software development and deployment pipeline so that every new deployment is automatically assessed for cryptographic impact:

- **Pre-commit / code review gates:** SAST rules that flag new introductions of quantum-vulnerable cryptographic API calls (e.g., direct RSA key generation, ECDH without hybrid wrapping, deprecated cipher suite configuration). These gates should not block development in early program phases but should generate warnings that

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

feed into the CBOM. As the program matures, consider making quantum-vulnerable introductions a build-break for Tier-1 systems.

- **Dependency scanning in build pipelines:** SCA tools that flag when a dependency update introduces or changes cryptographic library versions. When a library version change alters the set of available algorithms, this should trigger a CBOM review.
- **Container image scanning:** For containerized deployments, scan images for cryptographic libraries and configurations. Container orchestration platforms (Kubernetes) should have admission controllers or policy engines that can flag containers without compliant cryptographic configurations.
- **Infrastructure-as-Code (IaC) scanning:** For cloud-native organizations, scan Terraform, CloudFormation, Pulumi, and similar IaC templates for cryptographic configuration (TLS policies, encryption-at-rest settings, KMS key types). Block or flag IaC deployments that create resources with quantum-vulnerable-only encryption.
- **CBOM auto-generation on deployment:** Each production deployment should automatically generate or update CBOM entries for the deployed component. This can be achieved through integration between CI/CD metadata (component name, version, deployment target) and the CBOM system.

1.6.2 Passive Network Monitoring (Continuous)

Maintain ongoing passive monitoring of network traffic to capture cryptographic protocol negotiations in production:

- Deploy network taps or SPAN port monitoring at key network aggregation points (data center borders, cloud transit gateways, VPN concentrators, internet edge).
- Configure monitoring to capture TLS ClientHello/ServerHello messages, SSH key exchange, IPsec IKE negotiations, extracting negotiated cipher suites, protocol versions, certificate chains, and key sizes.
- Establish baselines for "normal" cryptographic behavior and alert on deviations: unexpected use of deprecated algorithms, new TLS endpoints appearing that were not in the CBOM, certificate chain changes, protocol downgrades.
- For organizations with dedicated cryptographic discovery platforms (Category 1 tools from Section 1.2), many of these platforms provide continuous passive monitoring as a built-in capability.

1.6.3 Change Management Integration

Make cryptographic impact assessment a standard component of the change management process:

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

- Add a "Cryptographic Impact" field to change advisory board (CAB) submission templates. Every change request should answer: "Does this change introduce, modify, or remove any cryptographic component? Does it affect algorithm selection, key material, certificate chain, or cryptographic library version?"
- For changes that do affect cryptographic components, require CBOM update as a condition of change approval and post-implementation review.
- Use change management tools (ServiceNow, Jira Service Management, etc.) to automatically tag cryptography-relevant changes and route them to the Inventory & Discovery workstream for CBOM reconciliation.
- Integrate inventory updates into the change management workflow so that asset tag changes, new deployments, decommissioning events, and configuration changes are all reflected in the cryptographic inventory promptly.

1.6.4 Scheduled Full-Estate Rescans

Even with CI/CD integration, passive monitoring, and change management integration, drift will occur. Schedule periodic full-estate rescans to catch what continuous mechanisms miss:

- **Quarterly:** Full network scan of all known subnets to detect new endpoints, changed configurations, and certificate expirations. Reconcile results against the CBOM; flag discrepancies for investigation.
- **Semi-annually:** Deep scan including code repository sweeps, cloud configuration audits, and HSM/KMS audit log reviews. This is particularly important for catching gradual drift in application-layer cryptography that passive network monitoring cannot see.
- **Annually:** Comprehensive re-assessment including manual investigation of the top critical systems, physical walkthroughs of OT/embedded environments, and third-party/vendor cryptographic re-verification.
- **Event-triggered:** Conduct unscheduled rescans after significant events: major system deployments, M&A integrations, data center migrations, vendor product upgrades, cryptographic vulnerability disclosures (e.g., a new CVE in OpenSSL or a NIST algorithm parameter change).

1.6.5 External Monitoring and Intelligence

Complement internal discovery with external monitoring:

- **Certificate Transparency (CT) log monitoring:** Continuously monitor CT logs for certificates issued against the organization's domains. This detects rogue or

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

unauthorized certificate issuance and reveals shadow IT services with TLS certificates.

- **External attack surface monitoring:** Use external attack surface management (EASM) tools to discover internet-facing endpoints and their TLS configurations from an adversary's perspective.
- **Cryptographic vulnerability intelligence:** Subscribe to cryptographic vulnerability feeds (NIST NVD, vendor security advisories, crypto-specific mailing lists) and automatically cross-reference new disclosures against the CBOM. When a vulnerability is disclosed in a cryptographic library, instantly determine which CBOM entries are affected.
- **Standards and algorithm monitoring:** Track NIST, IETF, ETSI, and national agency announcements for algorithm deprecation notices, parameter changes, and new standardization. When a standard changes (e.g., NIST deprecation timeline acceleration), automatically re-flag affected CBOM entries for Phase 3 re-scoring.

1.6.6 Alerting and Response Framework

Define a tiered alerting model for cryptographic posture changes:

Alert Level	Trigger	Response	Timeline
Critical	Newly discovered system using deprecated/broken crypto (e.g., TLS 1.0, RSA-1024) in production; unknown internet-facing TLS endpoint detected	Immediate triage; assess whether this is active exploitation risk; remediate or isolate within 72 hours	Same day
High	Quantum-vulnerable algorithm introduced in new Tier-1 deployment without documented justification; certificate expiry approaching without renewal plan	Workstream lead review; CBOM update; remediation plan within 2 weeks	Within 1 week
Medium	New cryptographic library version detected that changes available algorithms; CBOM drift detected in quarterly rescan	Investigate; update CBOM; assess impact on migration timeline	Within 1 month

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Informational	New non-critical system discovered with standard cryptographic configuration; routine certificate rotation completed	Log in CBOM; no action required unless pattern indicates broader issue	Next scheduled review

1.6.7 Organizational Culture and Training

Continuous discovery depends on organizational culture as much as tooling:

- Train developers and engineers to recognize and report cryptographic decisions in their work. Run awareness programs so that teams understand why documenting cryptographic choices matters and how to flag them through the established channels.
- Designate "crypto champions" in each major platform or application team who serve as the Inventory & Discovery workstream's point of contact and who proactively flag cryptographic changes within their team's domain.
- Establish feedback mechanisms (dedicated Slack/Teams channel, regular office hours, internal wiki) where anyone in the organization can report cryptographic observations that should be captured in the inventory.
- Include cryptographic inventory maintenance in performance objectives for the crypto champions and the Inventory & Discovery workstream lead. What gets measured gets done.

1.6.8 Measuring Discovery Effectiveness

Track the health of your continuous discovery capability with these operational metrics:

Metric	**Target**	**What It Indicates**
CBOM freshness	≥90% of entries updated within last 90 days	Whether the inventory is being maintained
Discovery-to-CBOM lag	New system appears in CBOM within 5 business days of deployment	Whether CI/CD integration is working
Drift rate	<5% discrepancy between CBOM and quarterly rescan results	Whether continuous mechanisms are catching changes

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Unknown-unknowns closure	Reduce gap register by 20% per quarter	Whether the organization is systematically closing discovery gaps
Time-to-detect	New unauthorized crypto endpoint detected within 7 days	Whether passive monitoring is effective

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

OUTPUTS

Output	Quality Criteria
Risk-driven scoping document	Tier-1 systems identified; discovery priorities A/B/C assigned; scoping rationale documented
Cryptographic asset inventory	Priority A systems: ≥90% coverage within 60 days; Priority B: ≥70% within 120 days; ≥90% of all Tier-1 systems within 12 months
Classical vulnerability findings	All classically vulnerable cryptography (deprecated algorithms, weak keys, expired certs, misconfigured protocols) identified and reported to security operations for immediate remediation
Sensitive data map (integrated)	All data classified by confidentiality requirement and retention period
Systems and assets register (integrated)	All assets classified by criticality and vendor dependency; cross-referenced with BIA classifications
Discovery gap register	Documented list of systems where discovery was incomplete, with remediation plan and target dates
Continuous discovery operating model	CI/CD integration operational; passive monitoring deployed; change management integration live; quarterly rescan schedule established; alerting framework defined; crypto champions designated

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

INTERDEPENDENCIES

Backward dependencies

See Prerequisites above. Phase 1 depends on Phase 0 outputs (charter, budget, governance, scoping assessment) and several organizational prerequisites including system access, baseline asset knowledge, and discovery tooling budget.

Feeds into

- **Phase 2** — Inventory data is the raw material from which the CBOM is populated. The data format and field completeness of discovery outputs directly determine CBOM quality, which is why the CBOM schema should be defined before large-scale discovery begins.
- **Phase 3** — The inventory is the primary input for risk scoring. System classification by business criticality (from Activity 1.0), exposure type, and data sensitivity provides the context needed to calculate priority scores.
- **Phase 7** — Discovery reveals vendor dependencies that may not have been visible during Phase 0 scoping. As each system is inventoried, the discovery team should flag vendor products that control cryptographic operations. This data feeds directly into the vendor classification matrix (Phase 7, Activity 7.1).
- **Phase 0 (feedback)** — Classical cryptographic vulnerabilities surfaced during quantum-focused discovery (deprecated protocols, weak keys, expired certificates) provide immediate near-term security wins that strengthen the business case and demonstrate program value to the executive sponsor and SteerCo.

Runs in parallel with

- **Phase 2** — CBOM structure should be defined in the first weeks of Phase 1 so that discovery data flows directly into a standardized format. Phase 1 and Phase 2 are tightly coupled and should be staffed as coordinated workstreams, not sequential handoffs.
- **Phase 7** — As discovery identifies vendor dependencies, preliminary vendor outreach (questionnaires, roadmap inquiries) should begin immediately for any vendor appearing on the critical path, particularly vendors whose products were not on the initial Phase 0 top-10 list.
- **Itself, indefinitely** — Discovery does not "complete." It transitions from project-mode initial discovery into a permanent continuous discovery capability. Staff and budget planning must reflect this.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

COMMON FAILURES

- **Interview-driven inventory.** Relying on application owners to self-report cryptographic usage. People don't know what crypto their systems use. Automated discovery supplemented by manual investigation is the only defensible approach.
- **Spreadsheet-only inventory.** A static spreadsheet becomes stale within weeks. The inventory must be maintained in a queryable system (CMDB, dedicated tool, or structured database) with automated refresh.
- **Waiting for 100% completeness before proceeding.** Completeness is asymptotic. You will never find everything. Use the 80/20 risk-driven scoping (Activity 1.0) to deliver actionable results quickly, proceed to Phase 2/3, and continue discovery in parallel.
- **Ignoring OT and embedded systems.** These are often the hardest to discover AND the hardest to migrate. Excluding them from scope guarantees surprises later. Physical walkthroughs are essential for OT environments.
- **Trusting a single tool for coverage.** No cryptographic discovery tool covers all five layers (network, code, config, runtime, manual). Deploy tools from multiple categories plus manual investigation.
- **CMDB-only asset discovery.** The CMDB is rarely complete. Organizations that rely solely on the CMDB for asset discovery consistently undercount their actual estate, particularly cloud resources, shadow IT, OT devices, and vendor-managed systems. Integrate the full range of data sources described in Activity 1.4.
- **Ignoring the immediate classical findings.** A quantum-focused cryptographic inventory will inevitably surface classically vulnerable cryptography (deprecated protocols, weak keys, expired certificates). Failing to remediate these findings immediately wastes a significant source of near-term security value and undermines the business case for the program. Report and remediate classical findings in parallel with continuing quantum-focused discovery.
- **Treating discovery as a project with an end date.** Discovery is not a phase that "completes"; it is a permanent operational capability. Organizations that disband the discovery team after "finishing" the initial inventory find their CBOM becomes stale within 6 months.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

MATURITY INDICATORS

Level	Indicator
Level 0	No cryptographic inventory exists; no awareness of the need
Level 1	Partial manual inventory of obvious systems (web servers, VPN); CMDB-only asset register; no continuous discovery
Level 2	Risk-driven scoping complete; automated discovery deployed on Priority A systems; ≥70% Tier-1 coverage; inventory is queryable; classical vulnerabilities being remediated; multiple asset data sources cross-referenced
Level 3	≥90% coverage; continuous discovery in CI/CD and passive monitoring; integrated with CMDB, SBOM, BIA, and certificate management; change management integration live; crypto champions designated; alerting framework operational
Level 4	Real-time cryptographic posture monitoring with tiered alerting; automated drift detection; coverage spans IT, OT, cloud, and third-party; discovery effectiveness metrics tracked and reported; discovery gap register trending toward zero

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PHASE 2 — CBOM & DOCUMENTATION

PURPOSE

Transform raw inventory data into a durable, queryable, standardized Cryptographic Bill of Materials (CBOM) as the single source of truth for all subsequent phases. The CBOM is the foundational artifact that makes PQC migration auditable, plannable, and measurable.

Without a structured CBOM, inventory data remains a collection of scan results, spreadsheets, and tribal knowledge scattered across teams. The CBOM transforms this raw material into a single, queryable, standardized record that answers the questions every subsequent phase asks: What cryptography does system X use? Is it quantum-vulnerable? What depends on it? Who owns it? What is its migration status? Organizations that skip or defer CBOM formalization find themselves re-discovering the same information repeatedly: every time a risk score needs calculating, a pilot needs scoping, or an auditor asks for evidence.

Parallelization note

Phase 2 should begin concurrently with Phase 1, not after it. The most common mistake is treating CBOM as something you build once discovery is "complete." In practice, the CBOM schema and tooling should be defined in the first weeks of Phase 1, so that discovery data flows directly into a structured format from the start. Phase 2 also runs in parallel with Phase 3: as CBOM entries are populated, risk scoring can begin on the entries that are ready, rather than waiting for full CBOM completion. The Minimum Viable CBOM model (below) is specifically designed to enable this parallelism: Layer 1 and Layer 2 data can be scored and prioritized while Layer 3 and Layer 4 discovery is still underway.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PREREQUISITES

Framework prerequisites (from earlier phases)

- **Phase 1 inventory data — at least partial**. The CBOM cannot be populated without discovery data, but it does not require a complete inventory to begin. As soon as Layer 1 (infrastructure) and Layer 2 (platform) discovery produces results, CBOM population should start. Waiting for Layer 3 and Layer 4 completeness before beginning CBOM work is a form of the completeness trap this framework warns against.
- **Phase 0 governance structure with designated data ownership**. Every CBOM entry needs an owner, someone accountable for the accuracy and currency of that record. The governance structure from Phase 0 should define how ownership is assigned (typically aligned to the system or service owner) and what "ownership" means in practice (update frequency, accuracy expectations, escalation for disputes).

Organizational prerequisites

- **A decision on CBOM format and hosting platform**. CycloneDX is the recommended standard (see Activity 2.1), but the organization must also decide where the CBOM will live: a dedicated tool, a CMDB extension, a version-controlled repository, or a purpose-built database. This decision has implications for integration with CI/CD pipelines, SBOM tooling, and reporting systems. Making this decision early avoids costly reformatting later.
- **SBOM maturity, or a plan to develop it in parallel**. The CBOM gains significant value when linked to Software Bill of Materials (SBOM) data, because SBOM reveals the dependency chains through which cryptographic libraries propagate. Organizations with no existing SBOM practice should plan to develop basic SBOM capability concurrently with CBOM, at least for Tier-1 applications.
- **CI/CD pipeline access (for organizations targeting automated CBOM updates).** If the organization intends to integrate CBOM generation into its software delivery pipeline, which is strongly recommended for Layer 3 coverage; the CBOM team needs access to CI/CD tooling and cooperation from development/DevOps teams. This access should be negotiated during Phase 0 or early Phase 1.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

THE MINIMUM VIABLE CBOM MODEL

The conventional approach to CBOM (attempting to catalog every cryptographic function call in every system before proceeding) is a completeness trap that delays migration indefinitely. The Minimum Viable CBOM (MV-CBOM) model takes an architecture-first approach organized in four layers:

- **Layer 1 — Infrastructure Cryptography**: TLS/SSH/IPsec configurations on load balancers, reverse proxies, VPN concentrators, and network devices. This layer is discoverable through network scanning and configuration review. It represents the largest attack surface for HNDL and is the most amenable to hybrid deployment.
- **Layer 2 — Platform Cryptography**: Cryptographic services provided by platforms, frameworks, and middleware (cloud KMS, HSMs, certificate authorities, identity providers, service mesh mutual TLS. This layer is discoverable through cloud API queries, HSM audit logs, and platform configuration review.
- **Layer 3 — Application Cryptography**: Cryptographic operations in application code, encryption of data at rest, digital signature generation/verification, token creation, custom protocol implementations. This layer requires code scanning and runtime analysis.
- **Layer 4 — Embedded/Third-Party Cryptography**: Cryptographic implementations in vendor products, firmware, IoT devices, and OT systems where the organization has no source code access and limited configuration control. This layer requires vendor documentation review, binary analysis, or acceptance of incomplete visibility.
- **The MV-CBOM strategy**: Achieve comprehensive coverage of Layers 1 and 2 first (weeks to months), because these layers contain the highest-exposure, most-controllable cryptographic usage. Achieve targeted coverage of Layer 3 for high-risk applications (months). Accept documented incompleteness at Layer 4, where vendor dependencies constrain visibility, and manage this through the vendor governance process (Phase 7).

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ACTIVITIES

2.1 Select CBOM Format and Tooling

Recommended format: CycloneDX, the de facto standard for CBOM, with native support for cryptographic asset types including algorithms, certificates, keys, protocols, and related dependencies. CycloneDX is supported by OWASP, adopted by the PQCC, and integrated into tooling from IBM, SandboxAQ, and the Linux Foundation's PQCA (CBOMkit).

CBOM record structure (per CycloneDX):

Each CBOM entry should capture at minimum:

Field	Purpose
Component identifier	Links to asset inventory and CMDB
Algorithm OID	Unambiguous algorithm identification
Key size / security parameter	Determines quantum vulnerability class
Protocol context	Where the algorithm is used (TLS, IPsec, S/MIME, etc.)
Implementation (library + version)	Identifies patching and upgrade path
Certificate reference	Links to certificate chain for PKI-related crypto
Data classification	From Track B of Phase 1
Quantum vulnerability status	Vulnerable (Shor), graded by quantum attack resource tier per algorithm and key size (see Phase 3); weakened (Grover); safe
Migration status	Not started / Planned / In progress / Hybrid / PQC-only / Not applicable
Owner	Responsible team for migration decisions
Vendor dependency flag	Whether migration is self-controlled or vendor-dependent

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

2.2 Populate CBOM from Inventory Data

Map Phase 1 inventory data into CBOM records. This is primarily a data transformation and enrichment exercise:

1. Import automated discovery results into CycloneDX format using tool-native exporters or CBOMkit
2. Enrich with manual investigation findings for systems not covered by automated tools
3. Cross-reference with SBOM data to establish library dependency chains
4. Link to certificate management system for PKI-related entries
5. Add data classification tags from Track B of Phase 1
6. Flag vendor dependencies from Track C of Phase 1

2.3 Integrate CBOM into Operational Processes

The CBOM has no value if it is a static document. Integrate it into:

- **CI/CD pipelines:** New deployments automatically generate or update CBOM entries. Block deployments that introduce quantum-vulnerable algorithms without documented justification and migration plan.
- **Change management:** CAB reviews include CBOM impact assessment: does this change introduce, modify, or remove cryptographic components?
- **Vendor onboarding:** New vendor products must provide CBOM-compatible cryptographic documentation as a procurement requirement.
- **Audit and compliance:** CBOM snapshots at regular intervals provide audit trail for regulatory evidence.

2.4 Establish CBOM Freshness Governance

Trigger	Action
New system deployment	Auto-generate CBOM entries via CI/CD
Library version update	Update implementation version; check for algorithm changes

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Certificate renewal/rotation	Update certificate reference and validity dates
Quarterly full scan	Reconcile CBOM against latest discovery results; flag discrepancies
Vendor product update	Request updated CBOM data from vendor; update entries
Algorithm deprecation notice	Flag all CBOM entries using deprecated algorithm; trigger Phase 3 re-scoring

2.5 Secure the CBOM and Program Artifacts

The CBOM is one of the most sensitive documents the organization will ever produce. A complete, current CBOM is a map of every weak cryptographic point in the enterprise: which systems still run quantum-vulnerable algorithms, where the unmigrated long tail sits, which vendor products block remediation, and which data stores carry the longest confidentiality requirements. To an HNDL adversary it is a collection shopping list; to any attacker it is target selection done for free. The same applies to the QRA and the prioritized migration backlog it produces.

This framework deliberately widens the CBOM's access surface (SIEM integration for SOC detection, GRC consumption for compliance reporting, CBOM exchange with vendors), which makes deliberate protection of the artifact non-optional:

- **Classify at the highest standard tier.** The CBOM, QRA, and migration backlog should carry the organization's most restrictive standard data classification (typically Restricted) and be handled accordingly.
- **Apply role-based access with least privilege.** Broad dashboard access should expose aggregated posture metrics, not queryable system-level vulnerability detail.
- **Log and review access.** Queries that enumerate unmigrated systems, or filter on quantum-vulnerable plus internet-facing, deserve the same scrutiny as queries against the vulnerability management database.
- **Control external sharing.** Vendors receive only the CBOM entries relevant to their own products; auditors and regulators receive point-in-time extracts through controlled channels, not standing access; consultant access ends with the engagement.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

- **Name CBOM exfiltration as a detection concern.** The SOC's data exfiltration monitoring (SOC Implementation, Use Case 5) should include the CBOM repository and GRC evidence stores among the high-sensitivity assets it watches.

No other element of the migration program concentrates comparable intelligence value in a single queryable asset. Protect it like one.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

OUTPUTS

Output	Quality Criteria
CycloneDX CBOM (Layers 1–2 complete)	100% of infrastructure and platform cryptography documented within 3 months
CycloneDX CBOM (Layer 3 targeted)	High-risk applications documented within 6 months
Layer 4 gap register	All vendor-dependent systems documented with known/unknown cryptographic usage
CI/CD integration	New deployments auto-generate CBOM entries
CBOM governance policy	Freshness rules, ownership, update triggers documented and enforced
CBOM protection controls	Restricted classification applied; role-based access and query logging operational; external sharing rules documented

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

INTERDEPENDENCIES

Backward dependencies

See Prerequisites above. Phase 2 depends on Phase 1 inventory data (at least partial), Phase 0 governance for data ownership, and several organizational prerequisites including CBOM format decisions and SBOM maturity.

Feeds into

- **Phase 3** — The CBOM is the primary data source for risk scoring. Each CBOM entry's algorithm, protocol context, data sensitivity, and migration feasibility fields are the inputs to the Phase 3 scoring model. CBOM entries that lack enriched metadata cannot be meaningfully scored, so CBOM quality directly constrains risk scoring quality.
- **Phase 5** — The CBOM's migration status field tracks pilot and production deployments, providing the authoritative record of what has been migrated, what is in progress, and what remains. Without this, migration progress reporting relies on manual tracking that becomes unreliable at scale.
- **Phase 7** — Layer 4 (embedded/third-party) CBOM gaps, where the organization cannot determine the cryptographic implementation because a vendor controls it, directly drive vendor engagement priorities. Each Layer 4 gap is a vendor question that Phase 7 must answer.
- **Phase 1 (feedback)** — CBOM analysis identifies discovery gaps: systems that should have CBOM entries but don't, dependencies that appear in one system's CBOM but whose source system hasn't been inventoried, or cryptographic usage patterns that suggest undiscovered systems. This feedback should drive targeted re-discovery.

Runs in parallel with

- **Phase 1** — CBOM structure should shape inventory data collection from the start. Discovery and CBOM population run as coordinated workstreams with a shared data model.
- **Phase 3** — As CBOM entries are enriched, risk scoring can begin on ready entries without waiting for full CBOM completion. This is particularly important for Layer 1 and Layer 2 entries, which should be scorable within weeks of initial discovery.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

COMMON FAILURES

- **Completeness trap.** Insisting on 100% CBOM coverage before proceeding to risk scoring and migration. Layer 4 (embedded/third-party) will never be fully visible. Accept this, document the gaps, and manage through vendor governance.
- **CBOM as document, not system.** Producing a CBOM in a PDF or spreadsheet that is never updated. The CBOM must be a live, queryable data set integrated into operational processes.
- **Ignoring the SBOM-CBOM linkage.** CBOM entries without SBOM context miss critical dependency chains. A vulnerable algorithm in a widely-shared library affects every application that depends on it.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

MATURITY INDICATORS

Level	Indicator
Level 0	No CBOM exists; cryptographic documentation is ad hoc or absent
Level 1	Partial CBOM in spreadsheet form covering known systems; no standard format
Level 2	CycloneDX CBOM operational for Layers 1–2; queryable; SBOM linkage established for key applications
Level 3	CBOM covers Layers 1–3; integrated into CI/CD; freshness governance enforced; change management integration live
Level 4	CBOM is a real-time operational asset; auto-updated on deployment; Layer 4 gaps systematically managed through vendor governance; CBOM drives automated compliance reporting

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PHASE 3 — RISK SCORING & PRIORITIZATION

PURPOSE

Translate CBOM data into a defensible, sequenced migration priority list. Not all quantum-vulnerable cryptography carries equal risk or equal migration difficulty. This phase produces the prioritized backlog that drives Phase 4 roadmap planning and Phase 5 execution sequencing.

Risk scoring is where the program transitions from "what do we have?" to "what do we do first?" Organizational politics often intrude here. Every business unit believes its systems are either the most critical (and therefore should be migrated first) or the least affected (and therefore should be left alone). A structured, transparent scoring model depoliticizes this conversation by replacing opinion with defensible, repeatable criteria. The QRA (Quantum Readiness Assessment) output from this phase is also the primary audit artifact: it demonstrates to regulators, auditors, and the board that the organization is making evidence-based decisions about migration sequencing, not simply reacting to whichever vendor or team shouts loudest.

Parallelization note

Phase 3 can begin as soon as the first CBOM entries are enriched with sufficient metadata. It does not require a complete CBOM. In practice, organizations should score Layer 1 and Layer 2 CBOM entries while Layer 3 discovery and CBOM population continue. This produces an initial prioritized backlog early enough to inform Phase 4 roadmap planning and Phase 7 vendor engagement. Risk scoring is also not a one-time activity. It must be re-run periodically (at least quarterly) as new inventory data arrives, regulatory deadlines shift, vendor PQC support timelines become clearer, and NIST standards evolve. The QRA should be treated as a living document, revised as these inputs change.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PREREQUISITES

Framework prerequisites (from earlier phases)

- **Phase 2 CBOM with enriched metadata — at least for Tier-1 systems**. Each CBOM entry being scored needs, at minimum: the algorithm(s) in use, the protocol context, the data sensitivity classification of the protected data, the system's business criticality tier, and an initial assessment of migration feasibility (can this be changed, and by whom?). CBOM entries that lack these fields cannot be meaningfully scored.
- **Phase 1 inventory data sufficient to identify Tier-1 and Tier-2 systems**. The risk scoring model requires a classified system inventory to weight business impact. If Phase 1 has not yet produced a classified list of systems by business criticality, this classification must be done as a preliminary step within Phase 3.
- **Phase 0 governance structure that defines who accepts risk**. Risk scoring inevitably produces results that some stakeholders will contest. The governance model must define who has authority to accept residual risk, override priority assignments, or defer migration, and the escalation path when disagreements arise. Without this, the scoring exercise produces a document that no one acts on.

Organizational prerequisites

- **An existing risk management framework (or willingness to adopt one for this purpose)**. The risk scoring model in this phase extends ISO/IEC 27005 and NIST SP 800-30 with quantum-specific dimensions. Organizations with a mature risk management practice should integrate the quantum scoring factors into their existing framework. Organizations without one will need to establish at least a basic risk scoring methodology, which is valuable well beyond PQC migration.
- **Data classification scheme**. The "Data Sensitivity" scoring dimension requires an organizational data classification standard (e.g., Public / Internal / Confidential / Restricted). If the organization does not have one, creating a fit-for-purpose classification for PQC risk scoring is an immediate prerequisite. This is another example of how PQC migration forces long-overdue security hygiene improvements.
- **Legal counsel engagement for long-retention data**. Organizations storing encrypted data with retention periods exceeding 10–15 years (medical records, financial archives, government records, legal documents) face compounding HNDL exposure. Legal counsel should be engaged to assess whether data retention policies, contractual confidentiality obligations, or regulatory requirements create additional urgency for specific data stores. This input directly affects the risk scoring weights.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ACTIVITIES

3.1 Define Risk Scoring Model

The risk scoring model below extends established information security risk management methodologies (ISO/IEC 27005, NIST SP 800-30, NIST RMF) with quantum-specific threat dimensions. Organizations that already operate a mature risk management framework should integrate these quantum-specific factors into their existing scoring methodology rather than creating a parallel system. The weights below represent a default starting configuration; organizations should calibrate weights based on their specific risk profile, sector, and regulatory context through a structured analysis involving both cryptographic and business stakeholders.

Score each CBOM entry across four dimensions:

Dimension 1 — Data Sensitivity × Exposure Window (HNDL Risk)

Factor	Weight	Scoring
Data confidentiality requirement (years)	High	>15yr = Critical; 10–15yr = High; 5–10yr = Medium; <5yr = Low
Data volume / flow rate	Medium	High-throughput channels score higher (more harvestable material)
Accessibility to adversary	High	Internet-facing = Critical; Partner-facing = High; Internal = Medium; Air-gapped = Low

Dimension 2 — Trust Infrastructure Criticality (TNFL Risk)

Factor	Weight	Scoring
Key/certificate lifetime	High	Root CA (20yr) = Critical; Intermediate CA (10yr) = High; End-entity (1yr) = Medium
Scope of trust dependency	High	Enterprise-wide PKI root = Critical; Single-app signing key = Medium
Signature verification period	Medium	Firmware signed for 15yr lifecycle = Critical; Session auth = Low

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Dimension 3 — Migration Feasibility

Factor	Weight	Scoring
Control posture	High	Full control (both endpoints) = Easy; Shared control = Medium; Vendor-dependent = Hard
Protocol/standard maturity	Medium	PQC-capable standard exists and is implemented = Easy; Standard exists but no vendor support = Medium; No standard yet = Hard
Interoperability constraints	Medium	Internal-only = Easy; Partner ecosystem = Medium; Public internet = Harder (but proven at scale)
System age and architecture	Medium	Modern microservices = Easy; Monolithic but maintained = Medium; Legacy unsupported = Hard

Dimension 4 — Regulatory and Compliance Pressure

Factor	Weight	Scoring
Specific regulatory deadline	High	Mandatory deadline within 2 years = Critical; Within 5 years = High; General expectation = Medium
Audit exposure	Medium	External audit scope = High; Internal audit scope = Medium; Not audited = Low
Contractual requirement	Medium	Customer/partner contract requires PQC = High; Expected soon = Medium; Not required = Low

Algorithm-Specific Vulnerability Weighting

Not all Shor-vulnerable cryptography is equally close to the threat. Quantum attack cost scales with key size, not with classical security strength, and this inverts a common intuition. ECC was adopted because 256-bit keys deliver the classical strength of 3,072-bit RSA; Shor's algorithm does not honor that bargain. At equivalent classical security (P-256 versus RSA-3072), breaking the ECC key requires roughly 2.6 times fewer logical qubits and two orders of magnitude fewer Toffoli gates. Even against RSA-2048 (a weaker classical target), current resource estimates place P-256 in the same striking

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

band, and ECDLP attack circuits have received far less optimization attention than RSA factoring circuits, which means the ECC estimates are the ones more likely to improve.

The scoring implication: at equal data sensitivity and exposure, do not deprioritize ECC-protected assets relative to RSA-protected ones. Estates that standardized on ECC for performance (which describes most modern infrastructure: TLS with X25519 or P-256 ECDHE, ECDSA certificates, secp256k1 in digital asset systems) carry at least as much quantum urgency as RSA-legacy estates. When calibrating Dimensions 1 and 2, treat ECC-256 and RSA-2048 as the same urgency tier, with RSA-3072 and above one tier behind. Record the algorithm and key size in each CBOM entry's quantum vulnerability field (Activity 2.1) so this weighting can be applied mechanically rather than re-debated per system.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

3.2 Calculate Priority Scores

For each CBOM entry, calculate a composite priority score:

Priority = (HNDL Risk × 0.35) + (TNFL Risk × 0.25) + (Regulatory Pressure × 0.25) + (Migration Feasibility Inverse × 0.15)

Note: Migration Feasibility is inverted: harder migrations may need earlier starts despite lower risk, because they require longer lead times. The weighting reflects that HNDL is the most immediately exploitable threat.

Decision Point — Urgency Classification:

Based on composite scores, classify each CBOM entry into migration tiers:

Tier	Criteria	Target Timeline
Tier 1 — Immediate	Critical HNDL exposure + high adversary accessibility + feasible migration (you control both endpoints)	Begin within 6 months; hybrid deployment within 12 months
Tier 2 — High Priority	High HNDL or TNFL risk + regulatory deadline within 3 years	Begin within 12 months; hybrid deployment within 24 months
Tier 3 — Standard	Medium risk + migration feasible with vendor coordination	Begin within 24 months; complete within regulatory deadline
Tier 4 — Long Tail	Lower risk OR migration currently infeasible (vendor dependency, no standard)	Monitor; compensating controls; migrate when feasible

3.3 Apply Migration Sequencing Logic

Sequencing operates within the Two-Track Migration Model defined in Phase 5: Track A (key exchange and confidentiality) and Track B (signatures, PKI, and authentication) run as parallel workstreams, not as a serial list. Earlier versions of this framework presented key exchange, then signatures, then data at rest as a single priority order; v2.1 retires that serialization, because treating signatures as second in line systematically delays the migrations with the longest lead times. Within each tier, sequence as follows:

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

1. **Track A — key exchange (TLS, VPN, key establishment protocols).** Sequence by exposure: internet-facing TLS and partner VPNs first, because each migrated channel immediately stops creating new HNDL-vulnerable sessions against traffic adversaries are harvesting now; then internal east-west traffic; then application-embedded key establishment.

2. **Track B — signatures and PKI (code and firmware signing, CA infrastructure, document signing).** Sequence by trust-anchor longevity and lead time: code and firmware signing first: a forged update cascades through every downstream system that trusts it, the SolarWinds attack pattern without the months of infiltration; then CA modernization and key-lifetime reduction, then end-entity and document signing. TNFL exploitation requires an operational CRQC, but signature migration carries the longest lead times in the program (PKI architecture decisions, toolchain support, HSM dependencies), which is exactly why Track B starts in parallel rather than waiting its turn.

3. **Data at rest (within Track A).** Re-encrypt long-lived archives with strong symmetric encryption (AES-256) or apply PQC-aware key wrapping (see Activity 5.5). Lower urgency than data in transit, because exploitation requires exfiltration in addition to future decryption capability, but stores with the longest confidentiality horizons must not fall off the roadmap.

3.4 Produce the Quantum Readiness Assessment (QRA)

Consolidate risk scoring into a formal Quantum Readiness Assessment document as the defensible basis for migration planning and regulatory evidence:

- **Executive summary:** Overall quantum risk posture; aggregate maturity score; comparison to regulatory deadlines
- **Heatmap:** Visual representation of risk across the estate, organized by system tier and domain
- **Prioritized migration backlog:** Sequenced list of all CBOM entries with tier assignments, target timelines, and owner assignments
- **Gap analysis:** Where current posture falls short of regulatory requirements, with specific remediation actions
- **Compliance mapping:** How QRA outputs map to NIST, CNSA 2.0, ETSI, sector-specific requirements

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

OUTPUTS

Output	**Quality Criteria**
Scored and tiered CBOM	Every CBOM entry has a priority score and tier assignment
Quantum Readiness Assessment (QRA)	Executive summary, heatmap, prioritized backlog, gap analysis, compliance mapping
Migration sequencing recommendation	Ordered backlog ready for Phase 4 roadmap planning

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

INTERDEPENDENCIES

Backward dependencies

See Prerequisites above. Phase 3 depends on Phase 2 CBOM data (at least for Tier-1 systems), Phase 1 inventory with system classification, Phase 0 governance for risk acceptance authority, and organizational prerequisites including a risk management framework and data classification scheme.

Feeds into

- **Phase 4** — The prioritized migration backlog is the primary input for roadmap planning. Tier assignments determine sequencing; priority scores determine resource allocation; migration feasibility ratings identify systems that require vendor support before they can be scheduled.
- **Phase 5** — Tier assignments determine pilot selection. The highest-priority, most-controllable systems from the Phase 3 backlog become the first pilot candidates.
- **Phase 7** — Risk scoring reveals which vendor dependencies are blocking Tier-1 migration. A CBOM entry scored as "Tier-1 priority, migration blocked by vendor" translates directly into an urgent vendor engagement action with executive escalation authority.
- **Phase 0 (feedback)** — The QRA is a primary audit artifact and board reporting input. It provides the evidence the executive sponsor needs to sustain multi-year funding by demonstrating structured, defensible decision-making. QRA findings may also trigger revisions to the business case. For example, if risk scoring reveals higher HNDL exposure than initially estimated.

Runs in parallel with

- **Phase 2** — Risk scoring begins as soon as CBOM entries are sufficiently enriched, without waiting for full CBOM completion. Scoring results also feed back into CBOM prioritization, helping Phase 2 teams focus enrichment effort on the entries that matter most.
- **Itself, iteratively** — Risk scoring is not a one-time exercise. The QRA must be re-run at least quarterly as new inventory data arrives, regulatory deadlines shift, vendor timelines change, and NIST standards evolve.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

COMMON FAILURES

- **Equal priority for everything.** Treating all quantum-vulnerable crypto as equally urgent. This leads to resource dilution and no meaningful progress on the highest-risk systems.
- **Ignoring migration feasibility in prioritization.** Assigning highest priority to systems that cannot be migrated yet (no vendor support, no standard) while ignoring high-risk systems that could be migrated today.
- **Static risk assessment.** Producing the QRA once and never updating it. Risk scores change as standards mature, vendor products ship, and regulatory deadlines approach.
- **Neglecting the legal and data-retention dimension.** Organizations storing long-lived encrypted data (medical records, financial archives, legal documents, government records) face a compounding legal risk: data encrypted today with quantum-vulnerable algorithms may become decryptable in the future, potentially violating data protection principles (such as GDPR's data security requirements) and contractual confidentiality obligations. The QRA should prompt a legal risk assessment alongside the technical one, evaluating whether data retention policies should be revised, whether additional protection layers (PQC key-wrapping, tokenization) should be applied to long-retention data, and whether data subjects or counterparties should be informed of the evolving risk profile. Engage legal counsel early; do not treat this as a purely technical exercise.
- **Assuming ECC's classical strength buys quantum time.** Migrating RSA first because it is classically weaker misreads the threat model. Quantum attack cost tracks key size: 256-bit ECC keys are at least as exposed as RSA-2048 and considerably more exposed than RSA-3072. Organizations that modernized from RSA to ECC reduced their classical risk and increased their relative quantum exposure. Apply the algorithm-specific weighting in Activity 3.1.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

MATURITY INDICATORS

Level	Indicator
Level 0	No quantum risk assessment exists
Level 1	Informal awareness of which systems are "probably vulnerable"; no structured scoring
Level 2	Formal risk scoring model applied to Tier-1 CBOM entries; prioritized migration backlog exists; QRA document produced
Level 3	QRA updated quarterly; all CBOM entries scored and tiered; migration sequencing drives Phase 5 execution; legal risk dimension assessed
Level 4	Continuous risk posture management; automated re-scoring when CBOM changes or regulatory deadlines shift; QRA integrated into enterprise risk register and audit cycle

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PHASE 4 — ROADMAP & GOVERNANCE

PURPOSE

Translate the prioritized migration backlog from Phase 3 into a multi-year execution plan with realistic milestones, resource allocations, vendor coordination timelines, and governance checkpoints. This phase establishes the PMO discipline to manage that complexity.

Phase 4 is where the program's ambition meets organizational reality. A prioritized backlog (Phase 3) tells you what needs to happen and in what order; a roadmap tells you when it will happen, who will do it, what it will cost, and what it depends on. Most PQC migration timelines are constrained by vendor readiness, hardware refresh cycles, regulatory deadlines, and the availability of scarce cryptographic engineering skills, not internal execution capacity. The roadmap must model these external constraints explicitly, not just internal work sequencing. Organizations that build PQC roadmaps as if they were building a software delivery plan, estimating internal effort and scheduling accordingly, discover that the actual critical path runs through vendor GA dates and hardware procurement lead times they did not account for.

Parallelization note

Phase 4 overlaps with Phases 5, 6, and 7 in practice. The roadmap is not a plan you complete and then hand off for execution. It is a living instrument that is updated as pilots reveal unexpected infrastructure requirements (Phase 6), vendor timelines slip or accelerate (Phase 7), and early migration waves generate lessons that change assumptions for later waves (Phase 5). In practice, the initial roadmap is drafted as soon as the first Phase 3 risk scoring outputs are available, and it is revised quarterly thereafter. Phase 4 also activates vendor engagement (Phase 7) immediately. Vendor engagement cannot wait for the roadmap to be "finished," because vendor lead times are typically the longest items on the critical path.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PREREQUISITES

Framework prerequisites (from earlier phases)

- **Phase 3 prioritized migration backlog**. The roadmap sequences work according to the tier assignments and priority scores from Phase 3. Without a scored backlog, roadmap construction devolves into stakeholder negotiation about whose systems go first, exactly the politicization that structured risk scoring is designed to prevent.
- **Phase 0 budget structure and multi-year commitment**. A roadmap without confirmed multi-year funding is aspirational fiction. The Phase 0 budget structure, including the phased funding model and benefit tracking approach — provides the financial framework within which the roadmap must fit.
- **Phase 1/2 data sufficient to estimate migration scope**. The roadmap must include realistic effort estimates, which require at least a rough understanding of what needs to change in each system. CBOM data for Tier-1 and Tier-2 systems should be available, along with enough inventory data to identify major vendor dependencies and infrastructure constraints.

Organizational prerequisites

- **PMO capability or access to program management professionals.** PQC migration at enterprise scale is a program, not a project. It requires program management discipline: dependency tracking across workstreams, resource management across business units, risk and issue management, and structured reporting. Organizations without existing PMO capability should plan to hire or contract program management resources specifically for this initiative.
- **Visibility into infrastructure refresh cycles, cloud migration plans, and vendor contract renewal dates**. The most cost-effective PQC migration strategy piggybacks on already-funded infrastructure changes. The roadmap team needs access to IT asset lifecycle data, cloud migration roadmaps, and procurement calendars to identify these integration opportunities. Without this visibility, the roadmap will propose standalone PQC upgrades that are more expensive and harder to fund than embedded ones.
- **Stakeholder alignment on the program's time horizon**. Phase 4 forces the organization to confront a 4–15 year execution timeline. If key stakeholders still believe PQC migration is a 12–18 month effort, the roadmap planning process will surface this gap in expectations. It is better to resolve this misalignment at the start of Phase 4 than to discover it when Year 2 budget requests are rejected.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ACTIVITIES

4.1 Define Year-1 Starter Plan (90-Day Governance Sprint)

The first 90 days set the organizational foundation. Deliver five things:

1. **Organization:** Executive sponsor and QRPM confirmed; SteerCo calendar and charter published; RACI matrix for all workstreams.
2. **People:** 10-to-20 person training cohort enrolled in PQC fundamentals; "crypto champions" designated per platform team.
3. **Plan:** Crypto-BOM v1 with ≥70% Tier-1 coverage complete; 2 hybrid pilots selected (TLS and VPN); vendor questionnaires sent to top 10 strategic vendors.
4. **Policy/Procurement:** Cryptographic policy updated (approved cipher suites include PQC hybrids; shorter key lifetimes mandated for new certificates); PQC and crypto-agility clause added to all new RFPs and contract renewals.
5. **KPI pack:** Baseline values established; Q+1 targets set; board reporting template ready.

Year 1 — Quarter-by-Quarter Detailed Plan:

Quarter	Governance & People	Discovery & CBOM	Pilots & Technical	Vendor & Policy
Q1	Appoint QRPM; convene SteerCo; publish charter and RACI; designate crypto champions; launch training cohort	Deploy Priority A discovery (internet-facing, Tier-1 systems); begin CBOM structure design; cross-reference with existing CMDB, BIA, and cert management data	Select 2 pilot targets (TLS + VPN); set up lab/staging environments; define pilot success criteria and SLOs	Send PQC roadmap questionnaires to top 10 vendors; draft updated procurement clause; begin cryptographic policy revision
Q2	First SteerCo milestone review; present initial findings to executive sponsor; refine	Complete Priority A discovery; begin Priority B; publish Crypto-BOM v1 (≥70% Tier-1 coverage); surface	Execute pilots in lab/staging; measure performance baselines; document findings	Collect vendor responses; classify vendors by PQC criticality; publish updated cryptographic

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

	RACI based on Q1 experience	and report classical vulnerabilities	and compatibility issues	policy and procurement language
Q3	Second SteerCo review; present QRA findings; request Year 2 budget refinement if needed; expand training to second cohort	Begin Phase 3 risk scoring on Priority A systems; populate CBOM migration status fields; integrate CBOM into CI/CD (initial)	Promote validated pilots to production canary (1–5% traffic); measure production SLOs; begin Wave 1 planning	Follow up with non-responsive vendors; evaluate bridging patterns for vendor-blocked systems; insert PQC clauses into upcoming contract renewals
Q4	Quarterly board report (first formal); present Year 1 achievements and Year 2 plan; finalize multi-year roadmap	Complete QRA for Tier-1 systems; Crypto-BOM v2 (≥90% Tier-1, ≥60% Tier-2); establish continuous discovery operating rhythm	Complete canary validation; begin Wave 2 (internal production); document validated migration patterns as repeatable playbooks	Vendor PQC scorecard published; escalation plan for non-compliant strategic vendors; Year 2 vendor engagement calendar set

4.2 Structure the Multi-Year Roadmap

Year	Focus	Key Milestones
Year 1	Foundation	Crypto-BOM v1; 2–4 hybrid pilots (TLS, VPN); vendor commitments secured; policy updates; team training
Year 2	Tier-1 Rollout	Hybrid/PQC deployed on all Tier-1 internet-facing endpoints; PKI key-lifetime reductions implemented; dual-sign pilots for code signing
Year 3	Bulk Migration	Tier-2 systems migrated; internal east-west traffic hybrid-enabled; HSM firmware upgrades; vendor product upgrades tracked

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Year 4	Long Tail & OT	Tier-3/4 systems addressed; OT gateway protections deployed; embedded device migration or containment; compensating controls where migration infeasible
Year 5	Hardening & Agility	Eliminate remaining compensating controls; achieve crypto-agility across the estate; transition from hybrid to PQC-only where ecosystem supports it; achieve target maturity level

4.3 Align to Infrastructure Refresh Cycles

Map migration tasks to existing planned expenditures:

Planned Refresh	**PQC Opportunity**
Data center network refresh	Deploy PQC-capable switches/load balancers; enable hybrid TLS at edge
SD-WAN / VPN upgrade	Require PQC support in new VPN concentrators; enable hybrid IPsec
Cloud migration wave	Configure PQC-capable TLS on cloud load balancers; use cloud-native PQC KMS
PKI renewal / CA migration	Deploy PQC-capable CA; issue hybrid/composite certificates; shorten key lifetimes
HSM replacement cycle	Procure PQC-capable HSMs (Thales Luna v7.9+, Utimaco Quantum Protect, or equivalent)
Vendor contract renewal	Insert PQC roadmap requirements, crypto-agility clauses, SLAs
Application modernization	Implement cryptographic abstraction layers; replace hardcoded algorithm calls with provider pattern

4.4 Establish PMO Structure for Scale

For programs exceeding 10,000 tasks, establish:

- **Work breakdown structure (WBS)** aligned to the 8 workstreams defined in Phase 0

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

- **Dependency mapping** between workstreams (see critical dependency chains below)
- **Critical path analysis** identifying the longest dependency chain. This determines the minimum program duration
- **Resource leveling** to avoid overloading shared teams (especially PKI, security architecture, and vendor management)
- **Risk register** with escalation triggers (e.g., "if vendor X misses PQC GA date by >6 months, escalate to SteerCo for alternate vendor evaluation")

Critical dependency chains (must be mapped and tracked):

Understanding these dependencies prevents the most common scheduling failures:

Upstream Activity	Must Complete Before	Why
CBOM v1 (Phase 2)	Risk scoring (Phase 3)	Cannot prioritize what you haven't inventoried
Risk scoring (Phase 3)	Pilot target selection (Phase 5)	Pilots must target the highest-priority systems, not the most convenient ones
HSM firmware upgrade (Phase 6)	Application PQC migration for HSM-dependent apps (Phase 5)	Applications that depend on HSM-protected keys cannot migrate until the HSM supports PQC key types
PKI CA modernization (Phase 6)	Dual-sign / PQC certificate deployment (Phase 5)	Cannot issue PQC certificates until the CA infrastructure supports PQC algorithms
Library upgrade (Phase 5)	Application hybrid enablement (Phase 5)	Applications cannot negotiate PQC if the underlying cryptographic library does not support it
Vendor PQC GA release (Phase 7)	Migration of vendor-dependent systems (Phase 5)	Cannot migrate systems where the vendor controls the cryptographic implementation until the vendor ships PQC support
Network middlebox testing (Phase 6)	Production hybrid deployment on affected network paths (Phase 5)	Middleboxes that cannot parse PQC handshakes will cause connection failures if not identified in advance

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Procurement policy update (Phase 0/4)	New system acquisitions (ongoing)	Without PQC in procurement requirements, every new acquisition potentially creates new quantum-vulnerable debt

The typical critical path for most enterprises is: Discovery → CBOM → Risk Scoring → Roadmap → Pilot → Production Rollout, with Vendor Engagement running in parallel from Q1 Year 1 and Vendor GA Releases gating Pilot and Rollout. The vendor dependency is usually the longest segment on the critical path, which is exactly why vendor engagement must start early, before the full CBOM is complete.

4.5 Manage the Roadmap as a Living Instrument

The roadmap is not a static Gantt chart that you create once in Phase 4 and follow mechanically. Quantum readiness is a multi-year program operating in a rapidly changing environment where standards, vendor timelines, regulatory deadlines, and threat intelligence all shift.

Quarterly roadmap review (standing SteerCo agenda item):

Present KPIs alongside leading indicators:

- **Quantum technology signals:** Logical qubit milestones, resource-estimate drops for breaking RSA/ECC, gate speed improvements, major lab announcements hitting roadmap milestones earlier than expected
- **Adversary signals:** Intelligence reports on HNDL campaigns (targeted interception and archival of encrypted data), supply-chain targeting of PKI/code-signing toolchains
- **Standards and regulatory signals:** NIST/IETF/ETSI/ISO algorithm updates, parameter changes, hybrid profiles moving to RFCs, new/earlier regulatory mandates
- **Vendor milestone signals:** Vendor GA dates met or missed, new PQC-capable products entering the market, FIPS validation completions

Decision playbook for when the roadmap needs adjustment:

- If KPIs are off track → Add resources, descope lower-risk items, pull forward bridging patterns
- If leading indicators worsen (e.g., CRQC timeline estimates shortened) → Move to accelerated track with pre-drafted resource request

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

- If performance issues surface at scale → Deploy offload/scale-out, adjust SLOs, extend pilot windows before broader rollout
- If a critical vendor misses a committed date → Activate the champion-challenger pattern (Phase 7); deploy bridging pattern; escalate contractually

4.6 Define Milestone Gates

Each phase transition requires a milestone gate review:

Gate	**Criteria**	**Decision Authority**
G0 → G1	Charter approved; budget committed; QRPM appointed	Executive Sponsor
G1 → G2	≥70% Tier-1 inventory complete; CBOM structure defined	SteerCo
G2 → G3	CBOM populated; risk scoring complete; QRA delivered	SteerCo
G3 → G4	Multi-year roadmap approved; Year 1 plan resourced	Executive Sponsor
G4 → G5	Pilot designs approved; success criteria defined; rollback plans documented	SteerCo
G5 → G6	Pilots validated; performance baselines established; Tier-1 rollout approved	SteerCo
G6 → G7	Infrastructure upgrades scheduled; vendor commitments tracked	QRPM

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

4.7 Pre-Draft the Accelerated Execution Profile

The decision playbook in Activity 4.5 says “move to accelerated track with pre-drafted resource request” if leading indicators worsen. That instruction only works if the accelerated track exists as a concrete artifact before it is needed. Organizations that pre-draft the profile will probably never invoke it; organizations that do not will improvise under the worst possible conditions.

The Accelerated Execution Profile is a pre-approved package with five elements. First, the trigger conditions: a credible reassessment of CRQC timelines by a source the SteerCo has designated in advance, a regulatory deadline compression affecting the organization’s jurisdictions, or a cryptanalytic event against a deployed algorithm. Second, the compressed sequence: discovery narrows to Priority A systems; the CBOM accepts documented Layer 3 and Layer 4 gaps; Waves 2 through 4 merge into a single accelerated rollout with relaxed canary thresholds; and lower-tier systems move directly to compensating controls (segmentation, overlay encryption, key-lifetime reduction) rather than full migration. Third, the pre-approved risk acceptances that the compressed sequence implies, signed by the risk owner in advance so that invocation does not wait on committee cycles. Fourth, the pre-drafted resource request: contractor capacity identified, emergency procurement clauses already negotiated with strategic vendors, and pre-authorized budget reallocation. Fifth, the activation authority: who can invoke the profile, and through which escalation path.

Exercise the profile annually as a tabletop scenario; the SOC Implementation foundation provides the exercise structure. A profile that has never been walked through will fail at the first step that depends on a person who has since changed roles.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

OUTPUTS

Output	Quality Criteria
Multi-year roadmap	5-year plan with annual milestones; aligned to refresh cycles; critical path identified
Year 1 detailed plan	Quarter-by-quarter with named owners, resource allocations, and success criteria
PMO operating model	WBS, dependency map, risk register, meeting cadence, escalation procedures
KPI dashboard	Baseline values set; targets for Q+1 defined; board reporting template operational
Policy updates	Cryptographic policy, procurement policy, change management policy updated with PQC requirements

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

INTERDEPENDENCIES

Backward dependencies

See Prerequisites above. Phase 4 depends on Phase 3 prioritized backlog, Phase 0 budget structure and multi-year commitment, Phase 1/2 data sufficient for scope estimation, and organizational prerequisites including PMO capability and visibility into infrastructure refresh cycles.

Feeds into

- **Phase 5** — The roadmap determines pilot selection, sequencing, resource allocation, and milestone gates. It specifies which systems are in Wave 1, what success criteria apply, and what the rollback triggers are.
- **Phase 6** — The infrastructure upgrade schedule (HSM replacement, PKI modernization, middlebox upgrades) is a roadmap deliverable. Phase 6 executes what Phase 4 plans and funds.
- **Phase 7** — The roadmap establishes the vendor engagement timeline: when each Strategic Blocking vendor must deliver PQC support to avoid becoming a critical-path blocker. This timeline gives vendor engagement its urgency and provides the basis for contractual commitments.

Runs in parallel with

- **Phases 5, 6, and 7** — The roadmap is not a plan that is "completed" and handed off. It is a living instrument, revised quarterly as pilot results (Phase 5) reveal unexpected complexity, infrastructure assessments (Phase 6) identify new constraints, and vendor timelines (Phase 7) shift. Roadmap governance (SteerCo reviews, milestone gates, contingency triggers) runs continuously alongside execution.
- **Phases 1 and 2** — Discovery and CBOM population continue throughout Phase 4 and beyond. New inventory data may reveal previously unknown systems that change the roadmap's scope or sequencing; the roadmap review process must accommodate these inputs.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

COMMON FAILURES

- **Planning in isolation from refresh cycles.** Building a PQC migration plan that ignores the organization's existing hardware refresh, cloud migration, and vendor contract renewal schedules. This results in "big-bang" budget requests that get rejected and missed opportunities to embed PQC into already-funded programs.
- **Underestimating the vendor dependency on the critical path.** The longest segment of most PQC migration critical paths is waiting for vendor products to ship PQC support. Organizations that start vendor engagement in Year 2 instead of Q1 Year 1 discover too late that their timeline is vendor-constrained and cannot be compressed with internal effort alone.
- **Single-track roadmap with no contingencies.** A roadmap that assumes every vendor delivers on time, every pilot succeeds first try, and no regulatory timeline changes. Build in explicit contingency triggers and pre-drafted acceleration/deceleration plans.
- **Treating the roadmap as a project plan rather than a program plan.** PQC migration is not a project with a defined end date. It is a permanent operational capability (crypto-agility). The roadmap should transition from "migration execution" to "ongoing cryptographic posture management" by Year 4–5, not simply declare victory and disband the team.
- **Governance without teeth.** Establishing a SteerCo that meets but does not make binding decisions, approve funding, or hold workstream leads accountable. The governance structure must have decision authority over budget, timelines, risk acceptance, and vendor escalation.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

MATURITY INDICATORS

Level	Indicator
Level 0	No migration plan exists
Level 1	Informal plan exists (spreadsheet, no governance); single-year horizon
Level 2	Multi-year roadmap approved; Year 1 plan resourced; SteerCo operational; KPI baseline set
Level 3	Quarterly roadmap reviews operational; dependency mapping maintained; refresh cycle alignment documented; vendor engagement tracked on dashboard
Level 4	Roadmap is a living instrument with quarterly updates; contingency triggers defined and tested; leading indicators monitored; program transitioning from migration execution to ongoing posture management

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PHASE 5 — PILOTS & MIGRATION EXECUTION

PURPOSE

Execute the migration through controlled pilots, validate patterns, and scale to production through waves. This phase operationalizes the roadmap from Phase 4, starting with the highest-priority, most-controllable systems and expanding systematically.

Phase 5 is where the program produces its first tangible security outcomes: systems that are actually protected against quantum cryptanalytic threats. It is also where the organization confronts the gap between theoretical migration design and production reality. Pilots will surface infrastructure incompatibilities, performance regressions, middlebox failures, and interoperability problems that no amount of planning can fully anticipate. This is expected and desirable. The purpose of the pilot phase is precisely to discover these problems in a controlled setting with tested rollback procedures, rather than at scale during production deployment.

Parallelization note

Phase 5 runs concurrently with Phase 6 (Infrastructure Modernization) and Phase 7 (Vendor Governance) in a tightly coupled feedback loop. Pilots reveal infrastructure bottlenecks that feed Phase 6 requirements; infrastructure upgrades enable broader pilot scope; vendor engagement in Phase 7 unblocks systems that depend on third-party PQC support. Organizations should not wait for infrastructure modernization to be "complete" before starting pilots. Early pilots on systems with existing infrastructure support generate lessons and organizational confidence that justify the infrastructure investment. Similarly, Phase 1 discovery and Phase 2 CBOM maintenance continue throughout Phase 5, as migration activities themselves generate updated cryptographic posture data that must flow back into the CBOM.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PREREQUISITES

Framework prerequisites (from earlier phases)

- **Phase 3 tier assignments identifying pilot candidates**. Pilot targets should come from the Tier-1 priority list, not from team convenience. The risk scoring output from Phase 3 determines which systems justify the investment and risk of early migration.
- **Phase 4 roadmap with resource allocation and sequencing for the first migration wave**. The roadmap provides the timeline, staffing plan, and milestone gates that govern pilot execution. Pilots launched without roadmap context tend to be underfunded, under-measured, and disconnected from the broader program.
- **Phase 2 CBOM entries for pilot candidate systems**. Each pilot target needs a detailed CBOM profile: which algorithms are in use, at which protocol layers, with what key sizes, through which libraries, and with what dependencies on external endpoints. Without this data, pilot design is guesswork.

Organizational prerequisites

- **Cryptographic engineering skills — internal or contracted**. Pilot design and execution require hands-on cryptographic engineering capability: configuring hybrid TLS, testing PQC key exchange, validating certificate chains, profiling performance impacts. This is a specialized skill set that most organizations do not have on staff. The skills gap assessment from the Cross-Cutting section should have been addressed by this point. If it hasn't, pilot quality will suffer and timelines will slip.
- **A test environment that mirrors production topology**. Pilots must be tested in environments that include the same middleboxes, load balancers, WAFs, monitoring tools, and partner endpoints as production. A pilot that succeeds in a clean lab environment but fails in production because of an undiscovered middlebox incompatibility is a wasted effort. If a suitable test environment does not exist, building one is a Phase 5 prerequisite that should be budgeted in Phase 4.
- **Documented rollback procedures and authority to execute them**. Before any pilot goes live, the team must have a tested, documented procedure for reverting to the pre-pilot cryptographic configuration, and the organizational authority to execute that rollback without committee approval during a production incident. This means pre-approved change management windows and pre-authorized rollback triggers.
- **Library and platform readiness for PQC algorithms**. The underlying cryptographic libraries (OpenSSL, BoringSSL, Bouncy Castle, platform-native crypto APIs) must support the target PQC algorithms at the required versions. Verify this before committing to pilot targets: discovering mid-pilot that a library upgrade is needed transforms a PQC pilot into a library migration project with different risks and timelines.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

DEPLOYMENT ENVIRONMENT CLASSIFICATION

Understanding Deployment Constraints

Before selecting pilot targets or designing hybrid deployments, organizations must classify their systems by deployment environment, because the environment determines when PQC can enter production, regardless of technical readiness.

The PQC standards are final. The cryptographic libraries support them. But for organizations operating in regulated environments that require FIPS-validated cryptography, there is a gap between "the algorithms are standardized" and "you can use them in production." This gap is a concrete constraint on every migration roadmap.

As of June 2026, no cryptographic module has achieved FIPS 140-3 validation with PQC algorithms within the approved boundary (though some validated modules, notably AWS-LC, include PQC operations as non-approved services). Multiple vendors have modules in the CMVP queue or at the Modules in Process (MIP) stage, but the CMVP process averages 18 or more months from submission to certificate. The earliest realistic FIPS 140-3 validated PQC software modules are expected in mid-2027. HSM modules with combined FIPS 140-3 Level 3 and PQC support will likely follow later. Meanwhile, FIPS 140-2 certificates move to the Historical list on September 21, 2026, creating a parallel migration pressure: organizations must transition to FIPS 140-3 validated modules even before PQC validation is available.

Environment Classes

Unrestricted environments have no regulatory or contractual requirement for FIPS-validated cryptography. These systems can deploy PQC today. Modern cryptographic libraries (OpenSSL 3.5 and later, Go 1.24 and later, Node.js 24.5 and later, BoringSSL, and all current major browser stacks) support NIST-standardized PQC algorithms by default. For these environments, the constraint is organizational readiness, not regulatory approval. Start pilots immediately.

FIPS-aware environments operate in contexts where FIPS validation is expected or preferred but not legally mandated, including many commercial deployments, enterprise IT systems not handling government data, and organizations whose policies reference FIPS without regulatory obligation. These systems can deploy PQC using CAVP-certified algorithms (the algorithm-level testing that precedes module-level validation). Document a risk acceptance decision for the interim period before full CMVP validation, noting that the algorithms themselves are NIST-standardized and CAVP-tested, and that full module validation is in progress industry-wide. This is a defensible interim posture.

FIPS-required environments are subject to regulation, contract, or policy that mandates FIPS 140-3 validated cryptographic modules: federal government systems, defense

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

contractor environments handling CUI, financial systems under certain regulatory frameworks, and healthcare systems subject to HIPAA technical safeguards. PQC cannot enter production in these environments until validated modules are available. Plan pilot and staging work now; target production deployment for the period following module validation (mid-2027 at the earliest for software modules). Use the interim period to validate patterns, train teams, update policies, and prepare infrastructure.

CNSA 2.0 / National Security System environments must comply with NSA CNSA 2.0 timelines: PQC required in new acquisitions from 2027, with full compliance deadlines of 2030 (software and firmware), 2033 (network equipment and operating systems), and 2035 (all remaining). These environments have the most demanding algorithm requirements (ML-KEM-1024, ML-DSA-87 at NIST Security Level 5) and the strictest validation requirements. The FIPS validation gap directly constrains the CNSA 2.0 execution timeline.

Roadmap Implications

The deployment environment classification feeds directly into Phase 4 roadmap construction. For each system on the migration backlog, the "earliest production deployment" date is constrained by the environment class. For unrestricted systems, deployment can begin now. For FIPS-aware systems, deployment can proceed with documented risk acceptance. For FIPS-required systems, plan for mid-2027 at the earliest. For CNSA 2.0 systems, the date is determined by the intersection of CMVP validation availability and CNSA 2.0 milestone deadlines.

Organizations operating across multiple environment classes (which includes most of the framework’s target audience) should sequence their migration to start with unrestricted and FIPS-aware environments, using these as pilots and proving grounds, while tracking CMVP validation progress for FIPS-required environments.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

TWO-TRACK MIGRATION MODEL

Confidentiality and Integrity: Two Parallel Migration Tracks

PQC migration addresses two distinct categories of quantum threat, each with different urgency drivers, deployment patterns, and infrastructure dependencies. v1.1 of this framework treated these as part of a single prioritized sequence. v2.0 makes this parallel structure explicit. Organizations that treat PQC migration as a single stream consistently under-invest in whichever track is not their starting point.

Track A: Confidentiality (Key Exchange) addresses Harvest Now, Decrypt Later. Adversaries are intercepting encrypted data today. The data they capture cannot be un-captured. Every day that key exchange remains quantum-vulnerable creates permanent exposure for data with long confidentiality requirements. The mitigation is hybrid key exchange -- deploy X25519+ML-KEM-768 on TLS, VPN, and IPsec channels carrying sensitive data. This track has the highest urgency: adversaries are harvesting encrypted traffic today.

Track B: Integrity and Authentication (Signatures and PKI) addresses Trust Now, Forge Later and the integrity of digital trust infrastructure. Digital signatures protect code integrity, firmware authenticity, certificate chain validity, document non-repudiation, and identity assertions. A quantum adversary with signature forgery capability could forge software updates, impersonate certificate authorities, and undermine the trust infrastructure that digital infrastructure depends on. Unlike HNDL, TNFL requires the adversary to possess a quantum computer at the time of the attack and cannot be executed retroactively. This means the urgency is lower than Track A, but the complexity and lead time are higher, because signature migration touches PKI architecture, code signing pipelines, certificate issuance infrastructure, and trust chain management.

Track B includes the identity stack, not only PKI. Most enterprise authentication rests on quantum-vulnerable signatures that a certificate-centric view misses. FIDO2/WebAuthn passkeys assert with ES256 or RS256 (P-256 and RSA, both Shor-vulnerable); IANA registered PQC COSE algorithm identifiers in April 2025 and IETF drafts for ML-DSA in WebAuthn are active, but deployable post-quantum passkeys remain years out. OAuth and OIDC tokens and SAML assertions are signed with RS256 or ES256. Kerberos PKINIT and smart card logon depend on RSA and ECC certificates, and 802.1X EAP-TLS authenticates devices with classical certificates. Two actions belong in scope now: inventory the identity stack as its own CBOM slice (identity providers, token signing keys and their rotation, authenticator algorithm policies, federation trust chains), and demand algorithm agility from identity vendors: every passkey enrolled today is a long-lived classical credential, and the question that separates real roadmaps from slideware is how existing credentials will be migrated or re-enrolled when signature algorithms

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

change. Treat short-lived session tokens as low priority (their TNFL exploitation window is bounded by token lifetime); treat long-lived credentials (passkeys, device certificates, token signing keys) as Track B assets under standard prioritization.

The two tracks also diverge because of the emergence of Merkle Tree Certificates as the likely architecture for post-quantum public Web PKI (see Phase 6). For public-facing PKI, Track B requires an architectural evolution, not merely an algorithm swap. Organizations need to plan for this now.

Why Two Tracks, Not One Prioritized List

Organizations that start with Track A (which most should) tend to declare progress when hybrid TLS is deployed on their top-priority endpoints. This is genuine progress -- but it addresses only confidentiality. If PKI modernization, code signing migration, and certificate lifetime reduction are not running as a parallel workstream, the organization accumulates a growing integrity gap that becomes harder to close as the program matures.

Conversely, organizations in sectors where TNFL is the primary threat -- particularly critical infrastructure and OT environments where compromised firmware signatures create safety risks, may need to start with Track B. These organizations still need Track A, but their risk profile demands parallel execution from the beginning.

Track Selection Guidance

For most enterprises, the recommended approach is: start Track A immediately with hybrid key exchange on Tier-1 data-in-transit channels. Launch Track B within 90 days: PKI lifetime assessment, certificate automation investment, code signing inventory. Maintain both as parallel workstreams with separate milestones and reporting.

Organizations handling classified data or operating national security systems should start both tracks simultaneously. CNSA 2.0 deadlines apply to both key exchange and digital signatures, and the PKI migration lead time for these environments is typically longer.

Organizations in critical infrastructure and OT should weight Track B higher than the default. TNFL is the primary quantum threat to safety-critical operations. Firmware signing, SCADA authentication, and control system integrity depend on signature infrastructure that may have 15-to-40-year equipment lifecycles.

Track B has a deploy-now component: stateful hash-based signatures. LMS and XMSS (NIST SP 800-208) are standardized, conservative, and required by CNSA 2.0 for software and firmware signing, the one signature migration NSA guidance says to begin immediately. Validated implementations exist, including in current HSM firmware. For firmware signing, software update signing, and secure boot chains (precisely the longest-lived trust anchors in the estate). The correct Track B action is not to wait for the ML-DSA

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ecosystem: deploy SP 800-208 signatures now, dual-signed with classical during transition. The operational constraint is state management: stateful schemes fail catastrophically if a key state is reused, so they belong in controlled, HSM-backed signing infrastructure, not general-purpose use. ML-DSA remains the path for high-volume, general-purpose signing as toolchains and validated modules mature.

Integration with Phase 5 Activities

The two-track model provides an organizing frame for the Phase 5 activity structure, not a replacement. Activity 5.1 (Select Pilot Targets) should produce pilot candidates for both tracks. Activity 5.2 (Design Hybrid Deployments) applies primarily to Track A. Activity 5.3 (Execute Pilots with Measurement) applies to both tracks with different metrics: latency and throughput for Track A, certificate lifecycle and signing pipeline compatibility for Track B. Activity 5.4 (Scale from Pilot to Production) applies to both with different wave structures. Activity 5.5 (Defense-in-Depth) supports both tracks.

PQC Deployment Ecosystem Status (June 2026)

Client-side PQC support is now the default across all major browsers and runtimes. The constraint has shifted to server-side and infrastructure readiness.

Browser support (PQC key exchange enabled by default): Chrome 131+, Firefox 132+ (desktop), Firefox 145+ (Android), Safari 26+ (including system-wide support in iOS 26 and macOS Tahoe 26), Edge 131+, Tor Browser 15.0+.

Library and runtime support (PQC enabled by default): OpenSSL 3.5.0+, Go 1.24+, Node.js 24.5.0+, BoringSSL (current), Rust (rustls 0.23.22+).

Adoption data: F5 Labs' mid-2025 measurement found 42% of the top 100 websites supporting hybrid PQC key exchange, dropping to 8.6% across the top one million, with only 3% of banking websites supporting PQC, and adoption has continued to climb since. Google and Cloudflare have both committed to completing their PQC migrations by 2029. Meta reports reaching its highest PQ-Enabled maturity level for portions of internal traffic.

For internet-facing services, the client side of most TLS connections already supports PQC without any action from the server operator. The business case for enabling server-side PQC is strengthened by the fact that clients are already requesting it.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ACTIVITIES

5.1 Select Pilot Targets

Choose 2–4 initial pilots based on:

- **Tier 1 priority** from Phase 3 risk scoring
- **Full organizational control** — you control both endpoints (no external partner dependency for initial pilots)
- **Measurable baseline** — current performance metrics exist for comparison
- **Representative architecture** — pilot results should be generalizable to other similar systems
- **Rollback capability** — must be able to revert to pre-pilot state if issues arise

Recommended starter pilots:

Pilot	Why First	What It Proves
Hybrid TLS on internet-facing web application	Highest HNDL exposure; well-understood protocol; strong library and browser support	PQC key exchange works at production scale; performance impact is acceptable
Hybrid IPsec/VPN between two controlled sites	High-value internal traffic; both endpoints controlled; clear performance baseline	PQC works for site-to-site encryption; validates VPN concentrator/firewall compatibility
Internal mTLS between microservices	Covers east-west traffic; high-volume; both endpoints controlled	Service mesh PQC capability; validates library compatibility
Certificate lifecycle with shortened validity	Not PQC algorithm change, but critical dependency for PQC readiness	PKI infrastructure can handle increased rotation frequency

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

5.2 Design Hybrid Deployments

Default recommendation for most enterprises: Hybrid ML-KEM-768 + X25519 for TLS key exchange as the initial deployment. In TLS 1.3, this is deployed as the X25519MLKEM768 named group. NIST SP 800-227 discusses the related X-Wing hybrid KEM construct as a general-purpose hybrid approach.

Hybrid cryptography runs classical and post-quantum algorithms together in the same operation. For a TLS handshake, this means combining X25519 (classical ECDH) with ML-KEM-768 (post-quantum key encapsulation). Both algorithms contribute to the shared secret, so the session remains secure as long as at least one algorithm is unbroken.

Why hybrid, not pure PQC:

1. Provides immediate quantum resistance without waiting for universal PQC adoption
2. Can preserve interoperability during phased rollout when the protocol and implementation support negotiating either classical or hybrid modes, though do not assume a hybrid deployment is automatically backward-compatible across every client, middlebox, or peer without testing
3. Reduces deployment risk — protected against both known quantum threats AND unknown weaknesses in new PQC algorithms
4. Proven at large production scale — Chrome/Chromium, Cloudflare, and AWS have all publicly documented production deployments of hybrid ML-KEM-based TLS, with Cloudflare reporting that a majority of human-generated TLS traffic to its network now uses PQC-hybrid key exchange

Hybrid signatures — this framework's position. The hybrid logic above is settled for key exchange; for signatures it is contested, and a vocal segment of practitioners argues for deploying ML-DSA alone to reduce complexity. This framework's position: composite or dual signatures (classical plus PQC) should be the default wherever toolchains support them, for the same defense-in-depth reason that made hybrid key exchange uncontroversial: new cryptographic code ships with implementation bugs, and ML-DSA implementations are new. Published analysis of early ML-DSA implementation flaws (Bernstein, June 2026) estimates that dropping the classical layer increases the share of breakable signature keys by roughly an order of magnitude over the next five years: a near-term implementation risk that exceeds the near-term quantum risk. The standard counterargument (that forgery damage is bounded by revocation) underestimates cascade effects in code signing and PKI, where a single forged signature propagates trust downstream before revocation takes effect. Where composite formats are not yet supported, dual-sign in parallel; where neither is feasible, deploying solo ML-DSA is an acceptance of measurable implementation risk and should be recorded as a risk decision, not treated as a default.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Hybrid deployment architecture per protocol:

Protocol	Classical Component	PQC Component	Standard/Draft	Production Ready?
TLS 1.3 Key Exchange	X25519	ML-KEM-768	IETF hybrid TLS drafts; deployed as X25519MLKEM768 named group; X-Wing (NIST SP 800-227) is a related general-purpose hybrid KEM construct	Yes — documented in production by Chrome, Cloudflare, AWS
IPsec IKEv2	ECDH P-256	ML-KEM-768	RFC 9370 (hybrid IKEv2)	Yes — available in major VPN products
SSH	X25519	ML-KEM-768	OpenSSH 9.9+ (`mlkem768x25519-sha256`; default from OpenSSH 10.0)	Yes — available in current OpenSSH
S/MIME / Email	ECDH P-256	ML-KEM-768	IETF drafts	Limited, not yet broadly deployed
Code Signing	ECDSA P-256	ML-DSA-65	Composite signatures (IETF)	Early — dual-sign approaches available
FIDO2 / WebAuthn	ES256 (P-256) / RS256	ML-DSA	COSE PQC identifiers (IANA, Apr 2025); IETF drafts	Not yet — standards groundwork only; require vendor algorithm agility
Token signing (JWT / OIDC / SAML)	RS256 / ES256	ML-DSA	JOSE/COSE ML-DSA drafts (IETF)	Early — inventory signing keys and rotation now

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Cryptographic library readiness: verify before committing to a deployment pattern:

The availability of PQC algorithms varies widely across cryptographic libraries. Before designing a pilot, verify that your specific library and version supports the algorithms you plan to deploy. As of early 2026, the major libraries include:

- **OpenSSL:** Native ML-KEM, ML-DSA, and SLH-DSA support available since version 3.5 (April 2025). Earlier versions required the external Open Quantum Safe (OQS) provider, which was experimental. Organizations still on OpenSSL 1.1.x face a major upgrade before PQC is available.
- **BoringSSL / Chrome:** ML-KEM implemented in September 2024 and enabled by default in Chrome 131, the fastest path to large-scale PQC deployment. BoringSSL's hybrid TLS support is the reference implementation for X25519MLKEM768.
- **AWS-LC:** First open-source library to achieve FIPS 140-3 validation with ML-KEM support. Critical for organizations requiring FIPS-validated PQC in AWS or federal environments.
- **Bouncy Castle (Java):** Supported NIST PQC finalists across versions released in 2022–2024, keeping pace with NIST drafts. Mature option for Java-based enterprise applications.
- **wolfSSL:** Full NIST PQC algorithm support aligned with CNSA 2.0 requirements. Strong position for embedded and IoT use cases.
- **Libsodium and MbedTLS:** As of early 2026, PQC integration is limited or still in progress. Projects dependent on these libraries may face delays and should evaluate alternatives or plan for wrappers.

The key implication: library readiness is a hard constraint on migration sequencing. If your application stack depends on a library that does not yet support PQC, the migration path is either (a) upgrade the library, (b) switch to an alternative library, or (c) deploy a PQC-capable gateway or proxy in front of the application as a bridge. Option (c) is often the fastest path for legacy applications.

Decision Point — Hybrid mandate divergence:

For multinational organizations, the hybrid approach debate creates compliance complexity:

- **BSI (Germany), ANSSI (France), Netherlands:** Strongly recommend or require hybrid PQ/traditional schemes (mandatory in some certification contexts)
- **NCSC UK:** Prefers a single migration to pure PQC over intermediate hybrid PKI (but accepts hybrid as interim)
- **NIST:** Permits but does not require hybrid

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

- **NSA CNSA 2.0:** Accepts hybrid as interim only

Recommendation for multinationals: Adopt hybrid as the default deployment pattern (satisfying the strictest requirement, continental European mandates) with the understanding that hybrid is an interim bridge to PQC-only as vendor and browser support matures. This "highest common denominator" strategy satisfies all jurisdictions simultaneously.

5.3 Execute Pilots with Measurement

For each pilot, define and measure:

Performance SLOs (Service Level Objectives):

- Handshake latency: Measure p50, p95, p99 before and after hybrid enablement
- CPU utilization: Server-side and client-side overhead delta
- Connection throughput: Impact on sustained data transfer
- Error rates: Connection failures, negotiation fallbacks, compatibility issues
- Memory utilization: Impact of larger key material and expanded state

Expected performance impacts (based on industry deployments):

- TLS 1.3 hybrid key exchange (e.g., X25519MLKEM768): expect a materially larger ClientHello and roughly ~1 KB of additional key-exchange data per handshake. The real latency impact is path-, MTU-, and middlebox-dependent; treat it as a measurement exercise. On modern server hardware, CPU overhead for ML-KEM operations is typically negligible; the dominant factor is packetization and middlebox behavior.
- IPsec hybrid: Similar latency impact; slightly higher CPU due to larger key operations
- Signature verification (ML-DSA): Signatures are far larger than ECDSA: ML-DSA-44 signatures are ~2,420 bytes and ML-DSA-65 signatures are ~3,309 bytes, versus ~64 bytes for ECDSA P-256. This is where real infrastructure stress appears.

Rollback criteria: Define explicit triggers for pilot rollback:

- Error rate exceeds X% above baseline
- Latency p99 exceeds Y ms above baseline
- Any security vulnerability identified in PQC implementation
- Interoperability failure with critical downstream system

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

5.4 Scale from Pilot to Production Through Waves

Wave	Scope	Prerequisite
Wave 0 — Lab/Staging	Isolated test environments only	Pilot design approved
Wave 1 — Internal, Non-Critical	Internal developer tools, monitoring systems, non-critical APIs	Lab validation passed
Wave 2 — Internal, Production	Internal Tier-1 services; east-west traffic	Wave 1 metrics acceptable
Wave 3 — External, Controlled	Partner-facing APIs with cooperative counterparties	Wave 2 stable; partner coordination complete
Wave 4 — External, Broad	Public-facing web applications, customer APIs	Wave 3 stable; external compatibility validated
Wave 5 — Long Tail	Legacy systems, OT gateways, embedded devices	Vendor support available or containment strategy deployed

5.5 Implement Defense-in-Depth Beyond Pure PQC

PQC algorithm migration is necessary but not sufficient. Complement hybrid/PQC deployment with:

- **Tokenization:** Replace sensitive data with tokens where possible, reducing the volume of data that requires cryptographic protection. Tokenization acts as a scope reducer for PQC migration: tokenized data stores eliminate entire categories of HNDL exposure.
- **Data minimization and retention reduction:** Reduce the data protection surface by deleting data that no longer has business value. Every data record you don't store is one you don't need to re-encrypt.
- **Network segmentation:** Limit adversary ability to harvest encrypted traffic by segmenting high-value data flows into isolated network zones with enhanced monitoring.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

- **AES-256 as universal symmetric default:** Ensure all symmetric encryption uses AES-256 (Grover-resistant) rather than AES-128. This is a low-effort, high-impact action that eliminates symmetric key vulnerability.
- **Key lifetime reduction:** Shorten certificate and key lifetimes where operationally feasible. A 90-day certificate limits the TNFL window compared to a 2-year certificate.
- **Confidential Computing as complementary protection:** For environments processing highly sensitive data, Confidential Computing (hardware-based Trusted Execution Environments) protects data in use, a gap that PQC alone does not address. PQC protects data in transit; AES-256 protects data at rest; Confidential Computing protects data during processing. Where available (e.g., Intel SGX/TDX, AMD SEV-SNP, ARM CCA), combining PQC transport security with Confidential Computing creates a defense-in-depth posture that is resilient against both quantum and classical platform-compromise threats.
- **Enhanced key wrapping for data at rest:** For archived data already encrypted with quantum-vulnerable key exchange, implement PQC-aware key-wrapping layers. Rather than re-encrypting entire data stores (which may be operationally infeasible for petabyte-scale archives), wrap the existing data encryption keys with PQC key encapsulation. This protects the keys without requiring data re-encryption, reducing the scope of the data-at-rest migration challenge.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

5.6 Decide the Data-at-Rest Strategy

Data at rest is deliberately sequenced after data in transit (Activity 3.3): an adversary must exfiltrate stored ciphertext before HNDL applies to it. But later is not never, and when the work begins it is not a single migration. It is a per-data-store choice among five strategies, driven by the confidentiality horizon of the data (from the Phase 1 classification), the volume, and the exfiltration exposure of the store.

Strategy	**When it fits**	**What to watch**
Re-encrypt under PQC-protected keys	High-sensitivity, long-horizon data in actively managed stores where re-encryption windows are operationally feasible	Throughput at scale; application downtime; verifying that no copy under the old keys survives
PQC key-wrap (KEK/DEK hierarchy redesign)	Large stores where re-encrypting the data itself is infeasible; archives with sound key hierarchies	Protects the keys, not ciphertext already exfiltrated; requires KMS/HSM PQC support; pairs with the double-wrapping pattern in Activity 5.5
Crypto-shred (destroy the keys)	Data past its confidentiality requirement that policy cannot yet delete outright	Only as strong as the certainty that no key copies exist; produce destruction evidence per the closure standard
Delete	Data with no remaining business or legal value	The cheapest migration is the data you no longer hold; requires legal sign-off on retention
Accept and monitor	Low-sensitivity or short-horizon data where the cost of any other strategy exceeds the risk	A recorded risk decision with a re-evaluation trigger, not a default

Backups and archives deserve their own line in the plan because they break the assumptions above. Tape libraries and immutable backups cannot be re-encrypted in place; the realistic pattern is to let old backup generations age out under a documented schedule while every new generation is written under PQC-protected keys, with key-wrap applied to long-retention archives that must outlive the transition. At petabyte scale, the decrypt-under-old, re-encrypt-under-new sequencing is a capacity-planning exercise

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

(Phase 6) as much as a cryptographic one. Whatever mix is chosen, record the strategy per data store in the CBOM, so that Phase 3 re-scoring and the closure verification in Migration Verification & Program Closure have an authoritative record of what was decided and why.

5.7 AI-Assisted Migration: Where It Helps, and the Gate That Stays Closed

AI tooling is now part of migration practice, and this framework takes a position rather than ignoring it. Where it genuinely accelerates: triaging discovery findings at Layer 3 scale, enriching CBOM entries from vendor documentation, generating test cases and interoperability matrices for migrated components, and drafting the repetitive engineering artifacts (configuration variants, rollout runbooks) that consume scarce cryptographic engineering hours. Fine-tuned models demonstrated roughly 92% functional correctness on cryptographic code migration in 2026 research. That figure argues for using the tooling on triage and drafting. It does not argue for trusting the output: at estate scale the remaining failures land somewhere, and in cryptographic code a functional-looking failure can be a security failure.

The non-negotiable gate: AI-modified cryptographic code passes through the same review, testing, and validation rigor as human-written cryptographic code, and where the change touches key generation, signing, or protocol state machines, more rigor, not less. This is the same evidence base that drives the composite-signature default in Activity 5.2: new cryptographic code ships with bugs, regardless of who or what wrote it. Algorithm selection, parameter choices, and risk acceptances remain human decisions. Record AI-assisted provenance on migrated components so reviewers and auditors know which artifacts carry it. The failure mode to refuse: treating AI throughput as a reason to compress review. The tooling buys speed in the steps before and after review, never in review itself.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

OUTPUTS

Output	Quality Criteria
Pilot results reports	Performance data (before/after), compatibility findings, issues log, recommendations
Validated migration patterns	Documented, repeatable patterns for TLS, VPN, mTLS, code signing
Wave deployment plan	Sequenced deployment schedule with success criteria per wave
Rollback procedures	Tested rollback capability for each deployment pattern
Defense-in-depth implementation	Tokenization, segmentation, AES-256 defaults deployed as complementary measures

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

INTERDEPENDENCIES

Backward dependencies

See Prerequisites above. Phase 5 depends on Phase 3 tier assignments, Phase 4 roadmap and resource plan, Phase 2 CBOM entries for pilot targets, and organizational prerequisites including cryptographic engineering skills, a production-mirror test environment, rollback authority, and library readiness.

Feeds into

- **Phase 6** — Pilot results are the primary input for infrastructure modernization scoping. Each pilot reveals which middleboxes fail, which HSMs need PQC key support, which network paths cannot handle hybrid handshake sizes, and which PKI components require modernization. Without pilot data, Phase 6 planning is based on assumptions rather than evidence.
- **Phase 7** — Pilot findings identify specific vendor gaps: a vendor product that fails hybrid negotiation, a middlebox that drops oversized ClientHello messages, a load balancer that cannot process PQC certificate chains. These findings provide the concrete evidence needed for vendor escalation, far more effective than abstract roadmap requests.
- **Phase 2 (feedback)** — Migration execution updates the CBOM. Each system migrated to hybrid or PQC changes its CBOM entry's migration status, algorithm fields, and compliance posture. Phase 5 is where the CBOM transitions from a discovery artifact to a live operational record.
- **Phase 4 (feedback)** — Pilot results (both successes and failures) generate lessons that change roadmap assumptions for subsequent waves. Actual effort data from Wave 1 should recalibrate effort estimates for Waves 2–N.

Runs in parallel with

- **Phase 6** — Infrastructure modernization and pilot execution are iterative. Pilots reveal infrastructure issues; infrastructure upgrades enable broader pilot scope. These two phases should be staffed and governed as a single tightly coupled workstream.
- **Phase 7** – Vendor engagement intensifies during Phase 5 as pilots generate specific, evidence-based requirements. The feedback loop is: pilot finds vendor gap → vendor engagement escalates → vendor delivers fix → next wave proceeds.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

- **Phase 6 dependency note**: Phase 5 both feeds into and depends on Phase 6. Early pilots can proceed on systems where existing infrastructure is adequate. Scaled deployment depends on infrastructure readiness. This circular dependency is resolved through the wave structure: Wave 1 pilots use existing infrastructure and generate Phase 6 requirements; later waves depend on Phase 6 delivery.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

COMMON FAILURES

- **Piloting the easy systems, not the important ones.** Choosing pilot targets based on convenience (the team that volunteers, the system with the friendliest owner) rather than Phase 3 risk prioritization. A pilot that succeeds on a developer sandbox proves little; pilots must prove that PQC works on the highest-priority, highest-risk systems.
- **Skipping rollback planning.** Deploying hybrid/PQC without a tested, documented rollback path. When a pilot causes unexpected failures, (and some will) the team must be able to revert within minutes, not hours or days. Define rollback triggers and test the rollback procedure before going live.
- **"Big bang" instead of waves.** Attempting to migrate all Tier-1 systems simultaneously rather than sequencing through controlled waves. Wave-based deployment (lab → internal non-critical → internal production → external controlled → external broad → long tail) allows each wave to validate assumptions and generate lessons for the next.
- **Ignoring library version as a hard constraint.** Designing a pilot for a system whose underlying cryptographic library does not yet support PQC algorithms. Verify library readiness (Section 5.2) before committing to a pilot target; otherwise the pilot becomes a library upgrade project, which has different risks and timelines.
- **Measuring only latency, ignoring compatibility.** Focusing pilot metrics on handshake latency while neglecting compatibility testing with downstream systems, middleboxes, monitoring tools, and partner endpoints. The most common production failures from PQC deployment are compatibility issues, not performance issues.
- **Treating hybrid as the end state.** Hybrid cryptography is a transition mechanism, not a permanent architecture. Organizations that deploy hybrid and declare victory will face a second migration later when hybrid is deprecated or when jurisdictions (NCSC UK, CNSA 2.0) require pure PQC. Plan the hybrid-to-PQC-only transition from the beginning.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

MATURITY INDICATORS

Level	Indicator
Level 0	No PQC pilots planned or underway
Level 1	Lab testing only; no production exposure
Level 2	2+ production pilots running with measured results; rollback procedures tested; validated patterns documented
Level 3	Tier-1 internet-facing systems on hybrid/PQC; wave rollout underway for Tier-2; defense-in-depth measures (tokenization, AES-256, segmentation) deployed
Level 4	Estate-wide hybrid/PQC deployment substantially complete; transitioning selected systems from hybrid to PQC-only; crypto-agility demonstrated via algorithm-swap drill

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PHASE 6 — INFRASTRUCTURE MODERNIZATION & PERFORMANCE

PURPOSE

Modernize the cryptographic infrastructure stack (PKI, HSMs, KMS, network devices, middleboxes) to support PQC operations at production scale. Address the performance, capacity, and compatibility challenges that PQC introduces to real infrastructure.

Phase 6 exists because PQC is not a "drop-in" algorithm swap. It changes the physical characteristics of cryptographic operations in ways that stress real infrastructure. ML-KEM ciphertexts are larger than ECDH key shares. ML-DSA signatures are dramatically larger than ECDSA signatures. Hybrid handshakes combine both, further increasing bandwidth and processing requirements. These changes propagate through every network path, every load balancer, every middlebox, every HSM, and every PKI component. Organizations that skip infrastructure assessment and testing discover these impacts in production, typically as intermittent failures that are difficult to diagnose because they manifest as network timeouts, dropped connections, or certificate validation errors rather than clean cryptographic failures.

Parallelization note

Phase 6 is not a discrete phase that happens after pilots. It runs iteratively alongside Phase 5. Early pilots reveal which infrastructure components need modernization; infrastructure upgrades enable broader and more ambitious pilots. PKI modernization in particular should begin early: reducing certificate lifetimes, testing dual-stack (hybrid) CA issuance, and inventorying HSM PQC capabilities are all activities that can and should start during Phase 1/2 timeframes, because they have long lead times and low risk. HSM procurement alone can take 6–12 months from order to deployment, and root CA

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ceremonies require extensive planning. Organizations that defer all Phase 6 work until Phase 5 pilots "prove the need" lose a year or more to infrastructure procurement and deployment lead times. Phase 6 also depends heavily on Phase 7: HSM vendors, PKI vendors, and network equipment vendors must deliver PQC-capable products on timelines that align with the migration roadmap, and vendor slippage directly delays infrastructure modernization.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PREREQUISITES

Framework prerequisites (from earlier phases)

- **Phase 5 pilot results identifying infrastructure bottlenecks and compatibility issues.** While some Phase 6 activities can begin proactively (HSM inventory, PKI lifetime reduction), the full scope of infrastructure modernization is informed by pilot findings. Pilots reveal which middleboxes fail, which HSMs lack PQC key types, and which network paths cannot handle hybrid handshake sizes.
- **Phase 4 infrastructure upgrade budget and scheduling.** Infrastructure modernization, particularly HSM replacement and PKI re-architecture, requires significant capital expenditure and long maintenance windows. These must be planned and funded in the Phase 4 roadmap.
- **Phase 1/2 inventory data covering infrastructure-layer cryptography**. The CBOM's Layer 1 (infrastructure) and Layer 2 (platform) entries identify which PKI components, HSMs, KMS instances, and network devices are in scope for modernization.

Organizational prerequisites

- **PKI team capacity and expertise**. PKI modernization is a specialized discipline. Organizations with outsourced or minimally staffed PKI operations will need to augment capacity before undertaking dual-stack CA deployment, root CA planning, or certificate lifetime reduction at scale. PKI errors cause outages. This work must be done carefully.
- **Network architecture documentation and middlebox inventory**. Testing PQC compatibility across network paths requires knowing what sits on those paths. If the organization does not have a current middlebox inventory (firewalls, IDS/IPS, WAFs, DLP, forward proxies, TLS inspection appliances), building one is a Phase 6 prerequisite. This inventory should include firmware versions and vendor support status, as middlebox PQC compatibility often depends on specific firmware releases.
- **HSM vendor engagement (from Phase 7).** HSM PQC capability is vendor-dependent. Before planning HSM modernization, the team needs vendor-confirmed information about which HSM models and firmware versions support PQC key types, and what the upgrade or replacement path looks like. This information comes from Phase 7 vendor engagement and should be requested in Q1 Year 1, well before Phase 6 begins in earnest.
- **Performance testing infrastructure**. Quantifying PQC's impact on handshake latency, throughput, and bandwidth requires load-testing capability that mirrors production traffic patterns. If the organization lacks a performance testing environment for cryptographic operations, establishing one is an early Phase 6 investment.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ACTIVITIES

6.1 PKI Modernization

PQC introduces fundamental challenges to PKI infrastructure:

- **Certificate size explosion:** ML-DSA-65 signatures are 3,309 bytes (ML-DSA-44 is 2,420 bytes) versus ~64 bytes for ECDSA. Certificate chains carrying PQC signatures are dramatically larger.
- **Increased chain verification time:** Larger signatures require more processing for chain validation.
- **Middlebox compatibility:** Many network devices (firewalls, IDS/IPS, WAFs, DLP) inspect TLS certificate chains. Larger certificates may exceed buffer sizes or parsing limits.

PKI migration strategy:

1. **Shorten key lifetimes immediately** — Reduce root CA key lifetimes from 20+ years to 10 years; intermediate CAs to 5 years; end-entity to 90–365 days. This limits TNFL exposure and builds organizational muscle for increased rotation frequency.
2. **Deploy dual-stack CA infrastructure** — Operate PQC-capable CA alongside existing classical CA. Issue hybrid/composite certificates for systems that need both. Maintain classical-only certificates for systems not yet ready.
3. **Track Merkle Tree Certificates (MTC) for web PKI** — Google and Cloudflare's MTC initiative (live-testing in Chrome as of early 2026 under the IETF PLANTS working group) addresses the certificate size problem structurally. Instead of sending full PQC signature chains over the wire, the CA signs a Merkle tree head, and the browser receives a compact inclusion proof (often less than 1 KB). This decouples cryptographic strength from bandwidth and integrates Certificate Transparency natively. The approach is still experimental and under active IETF standardization; it is the most promising path for PQC certificate size at web scale, but should not yet be treated as a near-term compliance obligation. Track IETF progress and plan for industry adoption.
4. **Test certificate chain processing** — Before any production deployment, test that all middleboxes (firewalls, WAFs, IDS/IPS, DLP systems, forward proxies) can process PQC certificate chains without failure or truncation.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

6.2 HSM and KMS Modernization

Action	Timeline	Consideration
Inventory all HSMs by model, firmware version, and PQC capability	Immediate	Many HSMs in production cannot support PQC without firmware or hardware upgrade
Upgrade HSM firmware to PQC-capable versions	As available	Thales Luna 7.8.0+ introduced initial PQC support (ML-KEM, ML-DSA); Luna 7.9 (June 2025) adds further PQC capabilities. Utimaco Quantum Protect is available as a new hardware variant (not a firmware upgrade for existing devices; budget for hardware replacement if running older Utimaco models)
Plan HSM hardware replacement where firmware upgrade insufficient	Align to refresh	Budget for 2–4 year replacement cycle
Configure cloud KMS for PQC key types	As available	AWS KMS, Azure Key Vault, Google Cloud KMS: check current PQC support status
Deploy software-based PQC key-wrapping overlay for HSMs not yet upgradeable	Bridge period	Use PQC key-wrapping around classical HSM operations as interim protection

6.3 Network Infrastructure Assessment

PQC impacts network infrastructure in several ways that must be tested before production deployment:

Handshake size impact:

- Classical TLS 1.3 handshake: ~1–2 KB
- Hybrid TLS 1.3 (X25519 + ML-KEM-768): ~2–3 KB (key exchange adds ~1 KB)

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

- PQC certificate chain with ML-DSA: Can exceed 10 KB for a 3-certificate chain

Protocol-specific concerns:

Protocol	PQC Challenge	Mitigation
TLS 1.3	Larger ClientHello may fragment; some middleboxes reject oversized handshakes	Test with production middleboxes; verify MTU handling; consider TLS certificate compression
DTLS (UDP)	Larger handshake fragments; amplification concerns	Test fragmentation handling; verify anti-amplification mechanisms
IKEv2/IPsec	Larger IKE_SA_INIT messages; fragmentation across VPN concentrators	Use IKE fragmentation (RFC 7383); test with production VPN devices
QUIC	Designed for larger handshakes; generally more PQC-friendly	Test with production load balancers
Constrained IoT (CoAP, LPWAN)	Extremely limited bandwidth; PQC key sizes may be prohibitive	Consider pre-shared key approaches; gateway-based PQC termination; sector-specific protocols
Satellite / LPWAN	High latency and limited bandwidth exacerbate handshake size impact	Use session resumption aggressively; consider PQC gateway at ground station
Mobile networks	Handshake latency affects user experience; bandwidth constraints on cellular	Test on representative cellular connections; measure impact on connection establishment time

6.4 Performance Testing Methodology

For each system targeted for PQC migration, execute:

1. **Baseline measurement:** Capture current performance metrics (latency p50/p95/p99, CPU utilization, memory, throughput) under representative production load
2. **Lab validation:** Deploy hybrid/PQC in isolated test environment; measure same metrics

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

3. **Canary deployment:** Deploy to 1–5% of production traffic with A/B comparison
4. **SLO evaluation:** Compare canary metrics against defined SLOs from Phase 5
5. **Capacity planning:** If PQC deployment requires additional compute/memory/bandwidth, estimate capacity requirements for full production rollout and include in Phase 4 budget

6.5 Capacity Planning for PQC at Scale

Resource	Expected Impact	Planning Action
CPU	ML-KEM key generation/encapsulation: minimal impact; ML-DSA signature verification: moderate	Benchmark on production hardware; plan for 5–15% server CPU increase for signature-heavy workloads
Memory	Larger key material and expanded TLS session state	Minimal for most systems; test on memory-constrained devices
Network bandwidth	Hybrid handshakes add ~1 KB; PQC certificates add 2–10 KB per connection	Significant for high-volume TLS termination points (CDN edges, load balancers); calculate aggregate bandwidth
Storage	Larger certificates and keys in certificate stores; larger OCSP responses	Manageable; plan for 2–5x increase in certificate storage

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

OUTPUTS

Output	Quality Criteria
PKI modernization plan	Dual-stack CA timeline; key lifetime reduction schedule; middlebox test results
HSM/KMS upgrade schedule	All HSMs inventoried with PQC capability status; upgrade/replacement timeline
Network compatibility report	All middleboxes tested with PQC handshakes; issues documented with mitigation
Performance baseline and projections	Before/after measurements for pilot systems; capacity plan for production rollout

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PKI ARCHITECTURE EVOLUTION

The PKI Architecture Fork

Post-quantum PKI migration is not a single path. As of mid-2026, two architectures are emerging, and organizations need to understand which applies to which parts of their infrastructure.

Path 1: X.509 certificates with PQC algorithms. The existing X.509 certificate chain architecture -- (root CAs, intermediate CAs, leaf certificates, Certificate Transparency logs) continues to work with post-quantum signature algorithms. Replace ECDSA or RSA signatures with ML-DSA. The trade-off is size: ML-DSA-65 signatures are approximately 3,309 bytes versus 64 bytes for ECDSA P-256. A typical TLS handshake carrying a certificate chain with five signatures and two public keys would see authentication overhead exceed 20,000 bytes -- roughly 30 times larger than today. For private and internal PKI -- enterprise mTLS, internal code signing, IoT device authentication, VPN certificate authentication -- this path remains viable and is the recommended approach.

Path 2: Merkle Tree Certificates (MTCs). Google announced in February 2026 a different architecture for public Web PKI, one that changes the trust model itself. Rather than signing each certificate individually with a large post-quantum signature, MTCs batch certificates into a Merkle tree and amortize a single post-quantum signature across thousands of certificates. The result: authentication overhead drops from approximately 14,700 bytes to as little as 736 bytes, making post-quantum HTTPS potentially smaller than today's classical certificate chains.

The MTC architecture is not a distant proposal. Chrome is implementing a Chrome Quantum-resistant Root Store (CQRS) targeted for Q3 2027. Let's Encrypt, which secures over 500 million websites, announced on June 3, 2026 that MTCs are its chosen path to post-quantum Web PKI, with staging environments in late 2026 and production issuance in 2027. Cloudflare is testing MTCs on live internet traffic. The IETF PLANTS working group has an active specification (draft-ietf-plants-merkle-tree-certs, published March 2026).

The Web PKI is heading toward MTCs. Organizations that operate public-facing web infrastructure should plan accordingly.

Planning Implications

Public-facing Web PKI (certificates for websites, public APIs, browser-authenticated endpoints): These deployments will transition to MTCs. Invest in ACME-based automated certificate management now -- the CA/Browser Forum Ballot SC-081v3 is reducing maximum certificate lifetimes to 200 days (March 2026), 100 days (March 2027), and 47

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

days (March 2029). Monitor Chrome and Let's Encrypt MTC implementation timelines. Engage your certificate authority about their MTC roadmap.

Internal enterprise PKI (mTLS, internal applications, VPN certificates, Wi-Fi/802.1X, smart card authentication): These deployments will use X.509 certificates with PQC algorithms. Begin PKI modernization activities now -- CA infrastructure PQC readiness, certificate lifecycle automation, HSM PQC capability assessment.

Code and firmware signing: These will continue using X.509 or equivalent structures with PQC algorithms. Implement dual-signing (classical + PQC) as a transition mechanism, plan signing key rotation to PQC, and assess toolchain support.

The Certificate Lifetime Convergence

The MTC transition intersects with the parallel move to dramatically shorter certificate lifetimes. The CA/Browser Forum Ballot SC-081v3 sets a trajectory from 398-day maximum to 47-day certificates by March 2029, with domain validation data reuse periods tightening to 10 days.

Organizations that invest in certificate lifecycle automation now -- regardless of PQC -- are simultaneously building the infrastructure they need for both the MTC transition and PQC PKI migration. Certificate automation is an infrastructure modernization that PQC makes urgent, not a PQC-specific activity. This is one of the strongest "benefit buckets" for the Phase 0 business case: the investment in automated certificate management pays for itself in operational efficiency, and the PQC readiness it enables is an additional benefit, not the sole justification.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

INTERDEPENDENCIES

Backward dependencies

See Prerequisites above. Phase 6 depends on Phase 5 pilot results, Phase 4 budget and scheduling, Phase 1/2 infrastructure-layer inventory data, and organizational prerequisites including PKI team capacity, middlebox inventory, HSM vendor engagement data from Phase 7, and performance testing infrastructure.

Feeds into

- **Phase 5** — Infrastructure readiness gates scaled production deployment. There is no point scheduling Wave 2 or Wave 3 migration if the PKI cannot issue hybrid certificates, the HSMs cannot store PQC keys, or middleboxes on the production path drop PQC traffic. Each Phase 6 deliverable (PKI dual-stack, HSM upgrade, middlebox clearance) unblocks a corresponding set of Phase 5 migration targets.
- **Phase 2 (feedback)** — Infrastructure modernization changes the CBOM. When an HSM is upgraded to support PQC key types, when a PKI root is re-issued with shortened lifetimes, when a middlebox firmware update enables PQC passthrough: these all update the infrastructure-layer CBOM entries and change the organization's overall cryptographic posture.

Runs in parallel with

- **Phase 5** — Infrastructure modernization and pilot execution are tightly coupled and iterative. Pilots reveal infrastructure issues; infrastructure upgrades enable broader pilot scope. These two phases form a continuous feedback loop.
- **Phase 7** — HSM, PKI vendor, and network equipment vendor timelines constrain infrastructure upgrade scheduling. Phase 6 execution depends on Phase 7 delivering vendor commitments and product availability. When a vendor misses a delivery milestone, Phase 6 must invoke a bridging strategy rather than simply delaying.

Early-start activities: Some Phase 6 activities should begin well before Phase 5 pilots produce results. HSM inventory (determining which models and firmware versions support PQC) should happen during Phase 1. PKI certificate lifetime reduction (a low-risk, high-value preparatory step) can begin as early as Phase 0/1. HSM procurement orders for models that need hardware replacement should be placed as soon as the inventory identifies the gap. Lead times of 6–12 months mean that deferring procurement until "Phase 6 officially starts" adds a year to the migration timeline.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

COMMON FAILURES

- **Assuming PQC is a "drop-in" library swap.** The most pervasive misconception. PQC changes key sizes, signature sizes, handshake characteristics, and certificate chain weights. Every network path, middlebox, load balancer, and TLS termination point must be tested, not just the application endpoints. Organizations that skip infrastructure testing discover in production that firewalls drop oversized ClientHello messages, IDS/IPS systems cannot parse PQC certificate chains, or WAF rules break on larger handshake payloads.
- **Ignoring middleboxes.** Firewalls, IDS/IPS, WAFs, DLP systems, forward proxies, and TLS inspection appliances sit on network paths between endpoints. Many have hardcoded buffer sizes or parsing assumptions that break with PQC-sized handshakes and certificates. Test every middlebox on every production network path. This is tedious but essential. Create a middlebox inventory as a sub-deliverable of Phase 6.
- **Deferring HSM upgrades.** HSMs are long-lead-time procurements with complex deployment procedures (key ceremony, compliance validation, high-availability configuration). Organizations that wait until Year 3 to discover their HSMs don't support PQC key types face 12–18 month delays for procurement, delivery, and deployment. Inventory HSM PQC capability in Year 1 and place orders immediately for any that need replacement.
- **Planning PKI modernization as a single event.** PKI changes (new CAs, shorter lifetimes, dual-stack operation) should be phased: start with end-entity certificate lifetime reduction (low risk, high value), then intermediate CA modernization, then root CA planning. Attempting all three simultaneously in a single maintenance window is a recipe for outages.
- **Neglecting capacity planning for high-volume TLS termination.** A CDN edge or load balancer handling millions of TLS handshakes per second will see measurable aggregate bandwidth increase from hybrid handshakes. Model the capacity impact before production deployment.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

MATURITY INDICATORS

Level	Indicator
Level 0	No awareness of PQC infrastructure impact; no testing
Level 1	Awareness of PKI/HSM/network challenges; no concrete plans
Level 2	HSMs inventoried with PQC status; PKI modernization plan drafted; initial middlebox testing underway
Level 3	HSM upgrades in progress; PKI dual-stack operational; all production middleboxes tested; performance baselines established for Tier-1 systems
Level 4	Infrastructure fully PQC-capable across IT estate; capacity planning validated at production scale; PKI automated with shortened lifetimes; middlebox monitoring integrated into continuous discovery

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PHASE 7 — VENDOR & SUPPLY CHAIN GOVERNANCE

PURPOSE

Ensure third-party products and services support the organization's PQC migration timeline. Vendor dependencies are the single largest external constraint on migration speed. "Our vendors will sort this out" is the most dangerous misconception in quantum readiness: vendors will update on their own timelines unless contractually compelled otherwise.

Phase 7 is numbered last in the framework's logical sequence, but in practice it is one of the first activities that must begin and one of the last to finish. Vendor PQC readiness timelines are typically the longest segment of the migration critical path, longer than internal development, longer than infrastructure procurement, longer than staffing ramp-up. An organization that executes Phases 0–6 flawlessly but defers vendor engagement until Year 3 will discover that it cannot migrate its most critical systems because the underlying vendor products do not yet support PQC. The vendor engagement process (questionnaires, roadmap alignment, contract negotiation, testing) also has its own lead time: 6–18 months from first inquiry to contractual commitment for strategic vendors. Every month of delayed vendor engagement is a month added to the end of the migration timeline.

Parallelization note

Vendor engagement should begin in Q1 Year 1, as soon as the top 10 critical vendor list is available from the Phase 0 scoping assessment. You do not need a finished CBOM or a complete risk scoring to send a PQC roadmap questionnaire to your most important vendors. Phase 7 then runs continuously throughout the program, escalating in intensity as CBOM data (Phase 2), risk scoring (Phase 3), and pilot findings (Phase 5) provide more specific and evidence-based requirements. Vendor governance does not end when migration "completes." It transitions into a permanent function that monitors vendor cryptographic posture, tracks algorithm deprecation timelines, and ensures ongoing supply chain quantum resilience. This is a permanent BAU function, not a project workstream.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PREREQUISITES

Framework prerequisites (from earlier phases)

- **Phase 0 initial scoping assessment, specifically the top 10 critical vendor list**. The scoping assessment should identify the vendors whose products sit on the critical path for Tier-1 and Tier-2 system migration. This list is the minimum input needed to begin vendor engagement; full CBOM and risk scoring data will refine and expand this list later, but waiting for that data before starting any vendor outreach wastes the program's scarcest resource: time.
- **Phase 1 vendor dependency data from inventory**. As discovery proceeds, the inventory reveals which systems depend on which vendor products for their cryptographic operations. This data feeds directly into the vendor classification matrix (Activity 7.1) and determines engagement urgency.
- **Phase 3 tier assignments determining vendor engagement urgency**. The risk scoring output tells you which vendor dependencies are blocking Tier-1 migration versus those affecting lower-priority systems. This determines where to apply executive-level escalation and contractual leverage versus standard engagement.

Organizational prerequisites

- **Procurement and legal team engagement**. Vendor governance for PQC requires contract modifications: adding PQC roadmap commitments, GA date requirements, testing obligations, and remedies for non-delivery. This requires procurement and legal involvement: security teams cannot unilaterally modify vendor contracts. Engage procurement and legal during Phase 0 and ensure they understand the program's vendor governance requirements before Phase 7 engagement begins in earnest.
- **Executive sponsor willing to escalate with vendor C-suites**. For Strategic Blocking vendors (those whose products are on the critical path and whose PQC timelines do not align with the organization's needs) engagement cannot remain at the technical or account management level. The executive sponsor must be prepared to escalate to vendor C-suite leadership, invoke contractual leverage, and if necessary authorize evaluation of alternative products. Without this executive backing, vendor engagement produces polite conversations and vague roadmap commitments that never materialize.
- **A structured vendor assessment methodology**. The organization needs a repeatable process for evaluating vendor PQC readiness: questionnaires, scoring criteria, tracking systems, and escalation triggers. Ad hoc inquiries produce inconsistent data and make it impossible to compare vendor readiness across the portfolio. The framework provides a classification matrix and questionnaire structure (Activities 7.1–7.2), but the organization must designate staff to operate this process.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

- **Awareness that open-source dependencies require the same governance discipline**. Many organizations focus vendor governance exclusively on commercial vendors while neglecting the open-source cryptographic libraries (OpenSSL, BoringSSL, Bouncy Castle, libsodium, language-specific crypto packages) that underpin their custom applications. These libraries have their own release cycles, PQC support timelines, and breaking change risks. The vendor governance process must include open-source dependency tracking with the same rigor applied to commercial vendors.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ACTIVITIES

7.1 Classify Vendor Portfolio by PQC Impact

Category	Criteria	Action
Strategic Blocking	Vendor product is in the critical path for Tier-1/2 migration; no PQC support = migration stops	Immediate engagement; executive-level escalation; contract leverage; parallel evaluation of alternatives
Strategic Enabling	Vendor product is PQC-relevant but alternatives exist or migration is not blocked	Standard engagement; require PQC roadmap; include in contract renewals
Non-Critical	Vendor product handles no quantum-vulnerable cryptography or protects only low-priority data	Monitor; include PQC in standard procurement language for renewals

7.2 Execute Vendor Engagement

For each strategic vendor, obtain:

1. **PQC feature matrix:** Which products support which PQC algorithms; current vs. planned
2. **GA version and dates:** When PQC support will be generally available (not beta, not "coming soon")
3. **Interoperability guidance:** Tested configurations with common counterparties
4. **Performance data:** Vendor's own benchmark data for PQC operation
5. **Fallback behavior:** What happens when a PQC negotiation fails (graceful degradation to classical, or hard failure?)
6. **CBOM data:** Cryptographic bill of materials for the vendor's product

Vendor PQC Readiness Questionnaire (core questions):

1. Which NIST-standardized PQC algorithms (ML-KEM, ML-DSA, SLH-DSA) does your product currently support?
2. For algorithms not yet supported: what is the GA date and version?
3. Does your product support hybrid/composite operation (classical + PQC simultaneously)?

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

4. Can cryptographic algorithms be changed via configuration without recompilation, reinstallation, or hardware replacement?
5. What is the performance impact of PQC operation (provide benchmark data)?
6. What testing have you done for interoperability with common ecosystem counterparties?
7. Will you provide CBOM data for your product in CycloneDX format?
8. What is your commitment timeline for FIPS 140-3 validated PQC implementations?

7.3 Insert PQC Requirements into Procurement

For all new purchases and contract renewals:

Add the following clause (adapt to legal counsel's preferred language):

"Supplier warrants that the Product/Service supports NIST-approved post-quantum cryptographic algorithms (FIPS 203, 204, and/or 205 as applicable) and provides crypto-agility (algorithm changes via configuration, not requiring recompilation or hardware replacement). Supplier will provide GA timelines, performance guidance, and interoperability matrices for PQC features. Failure to deliver PQC capability by [date] enables remedies up to and including [maintenance credits / right to replace / contract termination]."

For RFP requirements, add:

- Mandatory: NIST-standardized PQC algorithm support with published GA date
- Mandatory: Crypto-agility — algorithm selection via configuration
- Preferred: Hybrid operation support (classical + PQC)
- Preferred: CBOM data provided in CycloneDX format
- Required: FIPS 140-3 validation timeline for PQC implementations

Public model language now exists and is worth adapting rather than drafting from scratch. The Government of Canada's ITSM.00.501 procurement guidance and the Treasury Board's PQC requirements translate quantum readiness into dated, auditable contract terms; the U.S. DoD CIO memorandum on quantum-resistant cryptography does the same for defense suppliers; and CISA's guidance on procurement-relevant product categories signals where contract expectations are heading. Legal teams move faster when handed a government-grade exemplar.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

7.4 Manage Vendor-as-Blocker Scenarios

When a critical vendor cannot provide PQC support within your migration timeline:

1. **Deploy bridging patterns:** Software PQC overlay (gateway TLS termination, PQC key-wrapping around classical HSM, double-encryption layer) until native vendor support arrives.
2. **Run champion-challenger:** Evaluate an alternative vendor through POC while maintaining the incumbent. Communicate the POC to the incumbent vendor as leverage.
3. **Escalate contractually:** Invoke SLA remedies, maintenance credit provisions, or right-to-replace clauses.
4. **Accept and document risk:** If no alternative exists and no bridge is feasible, formally document the residual risk in the risk register with SteerCo approval and a re-evaluation trigger date.

7.5 Establish Ongoing Vendor Governance

Activity	Cadence
Track vendor PQC roadmap status	Monthly (for strategic blocking vendors); Quarterly (for others)
Verify vendor GA commitments against actuals	At each committed milestone date
Update vendor PQC scorecard	Quarterly
Report vendor status to SteerCo	Monthly
Re-evaluate vendor classification (strategic blocking / enabling / non-critical)	Semi-annually

Vendor KPI board (report to SteerCo):

- % of strategic vendors with signed PQC commitments and dates
- Number of products with GA hybrid/PQC capability in your estate
- Age of oldest unsupported critical product (days since vendor committed and missed)
- Number of bridging patterns deployed (indicates vendor gaps being compensated)

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

7.6 Coordinate Counterparties You Cannot Contractually Compel

Activities 7.1 through 7.5 assume leverage: contracts, SLAs, the right to replace. Partners, customers, API consumers, and interbank or B2B counterparties offer none of it, yet Wave 3 of the deployment sequence ("External, Controlled") depends on them moving. Counterparty coordination is negotiation, not governance, and it needs its own pattern. The pattern has six elements.

Readiness discovery. Determine each counterparty's actual PQC posture from observed protocol negotiation on shared connections (the passive monitoring from Activity 1.6 already captures it), supplemented by a lightweight readiness inquiry rather than the full vendor questionnaire.

Dual-stack windows. Commit to a published period during which both classical and hybrid negotiation are accepted on shared interfaces, so neither side is forced into a flag-day cutover.

Deprecation protocol. Announce classical-only retirement on a T-minus schedule (announcement, reminder at six months, enforcement date), in writing, through the commercial relationship owner as well as the technical channel.

Interoperability test events. Offer scheduled test windows against your hybrid-enabled endpoints before enforcement. Most counterparty failures surface here, where they are cheap.

API versioning. Expose PQC-capable endpoints as a parallel version rather than an in-place change, so counterparties migrate by choice of endpoint and their rollback is trivial.

The refusal decision. When a counterparty will not move, the options are a documented risk acceptance with compensating monitoring on that connection, commercial escalation through relationship owners, or, for connections carrying long-horizon data, termination. The decision belongs to the business relationship owner, made on the program's evidence, and it goes in the risk register either way.

The sector extensions operationalize this pattern for their networks (interbank messaging and card schemes in the Payments and Financial Services Extensions); the pattern above is the general case every sector needs.

7.7 Cloud Shared Responsibility and the SaaS Class

Cloud has appeared in this framework mainly as tooling surface (CSPM queries, KMS checks). The migration question is sharper: who migrates what. The provider migrates what it controls unilaterally (TLS termination on its managed front ends, the internals of managed services, its own inter-facility links), and the major providers are doing so on

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

their own schedules. The customer keeps everything above that line: application-layer cryptography, customer-managed keys and their KMS key types, workload-to-workload encryption, and every data-at-rest decision from Activity 5.6. The part most programs miss is boundary verification: do not mark a provider-side component migrated on the strength of a blog post. Capture provider attestations in the vendor register, and verify with observed algorithm negotiation at the service boundary wherever the connection is visible to you.

SaaS is a control-posture class of its own, not a variant of cloud. There is no bridging pattern for a SaaS application: you cannot terminate its TLS, wrap its keys, or proxy its internals. Vendor governance is the only lever. Classify each SaaS application in the CBOM by the sensitivity and confidentiality horizon of the data it holds, crossed with the vendor's PQC posture from the Activity 7.2 questionnaire. For high-sensitivity SaaS with a non-committal vendor, the realistic options are the Activity 7.6 refusal decision applied to a vendor you pay: risk-accept with monitoring, reduce the data entrusted (tokenization and minimization, Activity 5.5), or replace.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

OUTPUTS

Output	Quality Criteria
Vendor PQC readiness register	All strategic vendors assessed; GA dates recorded; blockers identified
Updated procurement requirements	PQC clauses in standard RFP templates and contract language
Vendor questionnaire responses	Documented responses from all strategic blocking vendors
Bridging pattern deployments	Deployed and documented for all vendor-blocked Tier-1/2 systems
Vendor governance cadence	Operating rhythm established with monthly/quarterly reviews

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

INTERDEPENDENCIES

Backward dependencies

See Prerequisites above. Phase 7 depends on Phase 0 scoping assessment (top 10 vendor list), Phase 1 vendor dependency data, Phase 3 tier assignments, and organizational prerequisites including procurement/legal engagement, executive escalation willingness, and a structured assessment methodology.

Feeds into

- **Phase 5** — Vendor GA releases gate production migration for vendor-dependent systems. A system cannot be migrated to PQC if the underlying vendor product does not yet support PQC algorithms. Each vendor GA delivery unblocks a set of Phase 5 migration targets.
- **Phase 6** — HSM, KMS, PKI, and network equipment vendor timelines constrain infrastructure upgrade scheduling. Phase 6 cannot upgrade an HSM that the vendor hasn't delivered PQC firmware for, or deploy a PKI platform whose vendor hasn't shipped PQC certificate support.
- **Phase 4 (feedback)** — Vendor timeline updates (both accelerations and delays) require roadmap revisions. A Strategic Blocking vendor that slips its GA date by 12 months changes the critical path for every system that depends on that product. The roadmap must model these vendor dependencies explicitly so that timeline changes propagate to affected milestones automatically.

Runs throughout the program: Vendor engagement is not a discrete phase with a start and end date. It begins in Q1 Year 1 with the initial top-10 vendor questionnaires and continues as a permanent governance function. As the program matures, vendor governance transitions from "getting commitments" to "verifying delivery," "testing interoperability," and eventually "monitoring ongoing vendor cryptographic posture" as part of business-as-usual supply chain risk management.

Escalating intensity over time: Early vendor engagement (Year 1) is primarily information-gathering: questionnaires, roadmap requests, classification. Mid-program engagement (Years 2–3) shifts to contractual commitment, testing, and escalation for non-delivery. Late-program engagement (Years 4+) focuses on verifying delivery, eliminating bridging patterns, and integrating vendor PQC support into standard procurement and renewal processes.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

COMMON FAILURES

- **Starting vendor engagement too late.** The vendor dependency is typically the longest segment of the PQC migration critical path. Organizations that defer formal vendor engagement until after the CBOM is "complete" lose 12–24 months. Begin vendor engagement in Q1 Year 1 using the top 10 vendor list from Phase 0 scoping. You do not need a finished CBOM to send a PQC roadmap questionnaire.
- **Accepting verbal commitments without contractual backing.** A vendor sales engineer saying "PQC is on our roadmap" is not a commitment. Insist on written GA dates, specific product versions, and contractual consequences for missed dates. Verbal assurances evaporate when vendor priorities shift.
- **Treating all vendors equally.** Vendors should be triaged by criticality (Strategic Blocking / Strategic Enabling / Non-Critical). Applying the same engagement intensity to all vendors wastes resources. Focus executive-level engagement and contractual leverage on the 5–10 Strategic Blocking vendors that constrain your migration timeline.
- **No bridging strategy for vendor gaps.** When a critical vendor cannot deliver PQC on your timeline, the response cannot be "wait and hope." Define bridging patterns (PQC gateway, overlay encryption, proxy-based TLS termination) proactively so they can be deployed immediately when a vendor misses a milestone.
- **The vendor delegation trap.** Assuming that because a vendor controls the cryptographic implementation, the vendor also owns the migration risk. The vendor owns their product; the organization owns the risk to its data, operations, and regulatory compliance. Vendor readiness is a dependency to manage, not a responsibility to delegate.
- **Ignoring open-source dependencies.** Many organizations track commercial vendor PQC readiness but neglect open-source components (OpenSSL, Bouncy Castle, language-specific crypto libraries) that underpin their custom applications. These require the same tracking discipline: monitor release schedules, test upgrades, and plan for version transitions.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

MATURITY INDICATORS

Level	Indicator
Level 0	No vendor engagement on PQC; assumption that "vendors will sort it out"
Level 1	Ad-hoc inquiries to a few vendors; no structured tracking
Level 2	Top 10 vendors formally engaged; questionnaires sent; responses tracked; vendor criticality classification complete
Level 3	PQC in standard procurement language; contracts include dated commitments and remedies; bridging patterns deployed for blocked systems; vendor scorecard reported to SteerCo
Level 4	All strategic vendors PQC-committed with verified delivery; bridging patterns eliminated as vendor support matures; open-source dependency tracking operational; vendor governance is permanent BAU function

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

MIGRATION VERIFICATION & PROGRAM CLOSURE

The eight phases define how to execute the migration. This section defines how to prove it happened and how the program ends. Both are routinely neglected: systems get reported as migrated on the strength of change tickets rather than observed behavior, and programs drift into an unfunded twilight instead of closing deliberately into a business-as-usual capability.

VERIFICATION: PROVING A SYSTEM IS MIGRATED

A migration status in the CBOM is a claim. Verification turns the claim into evidence. For each system marked Hybrid or PQC-only, apply the following evidence standard:

1. **Observed negotiation, not configuration.** Passive monitoring (the Phase 1 continuous discovery infrastructure) must show the system negotiating the expected hybrid or PQC algorithms in production traffic. Configuration states intent; handshakes state fact.

2. **Negative testing where classical-only must be refused.** For systems whose migration plan requires rejecting classical-only negotiation (post-transition environments, CNSA 2.0 scope), actively test that a classical-only client is refused. The absence of PQC failures does not prove the presence of PQC enforcement.

3. **Certificate chain validation under PQC.** For Track B migrations, verify the full lifecycle (issuance, validation by production clients, revocation) operates with PQC or composite certificates. Certificate issuance alone does not demonstrate this.

4. **Downgrade and fallback behavior documented.** Record what the system does when a counterparty cannot negotiate PQC, and confirm the behavior matches approved policy: graceful classical fallback during transition, refusal after.

5. **Evidence captured to the dossier.** Each verified system contributes a verification record (date, method, observed algorithms, tester) to the evidence dossier (see Metrics, KPIs & Reporting). This converts the dossier from self-reported progress into audit-grade proof.

Verification is sampled, not exhaustive, for large estates: verify 100% of Tier 1 systems and a documented sample (10–20% per migration wave) of lower tiers, with any sampled failure triggering full verification of that wave.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

DECOMMISSIONING CLASSICAL MATERIAL

Migration is not complete while the old keys remain. For each completed migration: revoke superseded certificates on schedule rather than letting them expire silently; destroy retired private keys under the organization's key destruction standard, with destruction certificates for HSM-held keys; remove retired roots and signing certificates from trust stores: a key is only retired when nothing trusts it; remove classical-only cipher suites from configurations once counterparty support permits; and update the CBOM entry to reflect the retired material. Retired-but-trusted classical signing keys are a standing TNFL liability: a forged signature from a retired key validates exactly as well as one from an active key if verifiers still trust it.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PROGRAM CLOSURE AND TRANSITION TO BAU

Define closure criteria in advance: typically maturity Level 4 or higher across all seven domains; verification complete for Tier 1 and Tier 2 systems; and every remaining item carried as an explicit risk acceptance with an owner and a re-evaluation date.

At closure, execute a formal handover. The QRPM role transitions to a standing cryptographic governance owner. The CBOM and continuous discovery, the vendor governance cadence, SOC detection content, KRI reporting, and crypto-agility OKRs transfer to named permanent owners with funded budgets. The evidence dossier is archived to the records system under the retention the organization's regulators require. The SteerCo formally accepts the residual risk register, and the closure decision is minuted at the level that chartered the program: the same governance that opened the program closes it.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

OUTPUTS

Output	Quality Criteria
Verification records	Evidence standard applied; 100% Tier 1 coverage, documented sampling for lower tiers; records filed in the evidence dossier
Decommissioning log	Certificates revoked on schedule; keys destroyed with certificates; trust stores cleaned; CBOM entries updated
Closure decision record	Closure criteria met or explicitly risk-accepted; decision minuted by the governance level that chartered the program
BAU handover register	Named permanent owner, funded budget, and operating cadence for every continuing capability

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

COMMON FAILURES

- **Declared done.** Closing on milestone completion rather than verified posture. If verification evidence does not exist, the migration is a belief, not a fact, and the first auditor, regulator, or counterparty who asks for proof will treat it accordingly.
- **Orphaned capabilities.** Treating closure as the end of funding rather than a handover. The CBOM, discovery infrastructure, and vendor governance decay within quarters if they close with the program instead of transferring to funded owners.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

PROGRAM FOUNDATIONS: CAPABILITIES THAT SPAN EVERY PHASE

The eight phases above describe *what* the program does and in *what order*. But several capabilities do not belong to any single phase. They operate continuously across the entire program lifecycle, shaping how every phase is executed, measured, and governed. These foundations are not optional supporting material. A program that executes the phases without building these capabilities will produce a migration that cannot be measured, cannot be sustained, cannot be staffed, and cannot satisfy regulators.

This section covers five foundational capabilities: a maturity model for benchmarking progress, metrics and KPIs for tracking and reporting, crypto-agility as the program's architectural end-state, regulatory and standards alignment for compliance mapping, and the skills and team structures needed to execute the work. Each should be established early, ideally during Phase 0 and Phase 1, and maintained throughout the program's lifetime.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

MATURITY LEVELS

Level	Name	Description
0	**Unaware**	No organizational awareness of quantum cryptographic risk; no activities planned or underway
1	**Aware**	Quantum risk acknowledged at leadership level; initial education conducted; no formal program
2	**Initiated**	Formal program chartered; budget allocated; discovery underway; initial CBOM in development
3	**Progressing**	CBOM operational; risk scoring complete; hybrid pilots in production; vendor engagement active; KPIs reported
4	**Advanced**	Tier-1 systems migrated to hybrid/PQC; PKI modernized; crypto-agility demonstrated; vendor commitments secured
5	**Optimized**	Estate-wide PQC migration substantially complete; crypto-agility is organizational capability; continuous monitoring and posture management operational

Assessment Across Seven Domains

Domain	Level 1	Level 2	Level 3	Level 4	Level 5
Cryptographic Inventory	Awareness that inventory is needed	Partial inventory of known systems	≥70% Tier-1 coverage; automated discovery deployed	≥90% coverage; continuous monitoring; IT+OT+cloud	Real-time posture management; auto-alerting on drift
Governance & Ownership	Informal awareness	Charter exists; QRPM appointed	SteerCo operational; RACI defined; budget committed	Multi-year governance; integrated into risk register	Quantum readiness part of BAU enterprise governance

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Pilots & Deployment	No pilots	Lab testing only	2+ production pilots with measured results	Tier-1 systems on hybrid/PQC; wave rollout underway	Estate-wide deployment; transitioning to PQC-only
Vendor & Supply Chain	No vendor engagement	Ad-hoc inquiries	Top 10 vendors formally engaged; questionnaires sent	PQC in procurement; contracts include clauses; blockers managed	All strategic vendors PQC-committed; bridging patterns eliminated
Compliance & Standards	Unaware of requirements	Regulatory requirements mapped	Compliance gaps identified; remediation planned	Meeting current deadlines; evidence documented	Proactive compliance; contributing to standards development
Crypto-Agility	Hardcoded algorithms throughout	Awareness of agility need	Abstraction layers in new development	Algorithm swap via config for Tier-1 systems; automated cert lifecycle	Organization-wide agility; algorithm changes are routine operations
Risk & Prioritization	No quantum risk assessment	Basic awareness of HNDL/TNFL	Formal risk scoring; prioritized backlog	QRA updated quarterly; migration tracking against backlog	Continuous risk posture management; automated re-scoring

Self-Assessment Scoring

For each domain, score 0–5. Overall maturity level is the lowest individual domain score (the weakest domain constrains overall readiness).

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Target maturity timeline:

- End of Year 1: Level 2 across all domains
- End of Year 2: Level 3 across all domains
- End of Year 3: Level 4 across at least 5 of 7 domains
- End of Year 5: Level 4 or 5 across all domains

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

METRICS, KPIS & REPORTING

Board-Level KPI Pack (Report Quarterly)

KPI	What It Measures	Target
Coverage	% of Tier-1 endpoints using hybrid/PQC key exchange	Year 1: 10%; Year 2: 60%; Year 3: 95%
Trust	# of PKI/signing anchors with shortened lifetimes and/or PQC capability	Year 1: Root CAs assessed; Year 2: Intermediates shortened; Year 3: End-entity automated
Inventory	% of RSA/ECC locations mapped and risk-ranked in CBOM	Year 1: 70% Tier-1; Year 2: 90% all tiers; Year 3: continuous
Vendors	% of strategic suppliers with dated PQC commitments	Year 1: 50%; Year 2: 80%; Year 3: 95%
Agility	% of services that can swap key-agreement algorithm via configuration within 2 weeks	Year 1: baseline; Year 2: 30%; Year 3: 60%

Risk-weighted coverage (companion metric). The Coverage KPI counts endpoints equally, which makes 60% read as progress even when the unprotected 40% carries the highest-value flows. Report an exposure-weighted companion: the share of HNDL-relevant traffic volume, or of long-retention data flows, protected by hybrid or PQC. The inputs already exist: traffic volumes from the passive monitoring in Activity 1.6, confidentiality horizons from the Phase 1 classification. Offer it alongside the endpoint metric, not as a replacement; the pair answers the question a regulator will actually ask, which is not how many systems, but how much of what matters.

Operational KPIs (Report Monthly to SteerCo)

KPI	Measurement
Migration tasks completed vs. plan	Earned value / planned value (by workstream)

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Pilot performance delta	p95 latency increase vs. baseline (target: <2ms for TLS)
Discovery gap closure	% of Phase 1 gap register items resolved
Vendor milestone adherence	% of vendor GA commitments met on time
Policy compliance	% of new deployments passing PQC CI/CD check
Training coverage	% of target staff completing PQC training modules

Evidence Dossier (for Audit and Regulatory)

Maintain a continuously updated evidence package:

- CBOM snapshots (quarterly)
- Pilot test reports with SLO data
- Vendor questionnaire responses and commitment letters
- Updated cryptographic and procurement policies
- SteerCo meeting minutes and decision records
- Training completion records
- Risk register with quantum risk entry and mitigation status

The dossier as litigation defense. Beyond audit and regulatory response, the dossier serves a purpose most programs only discover when it is too late to create one: negligence defense. When harvested data is decrypted years from now, or a forged-signature incident lands, the legal question will be whether the organization exercised reasonable care against a risk that was foreseeable and well documented at the time. The dossier is the answer, provided it captured three artifact types deliberately: risk-acceptance decisions with their rationale and re-evaluation dates; budget requests, including the ones that were denied (a logged, declined funding request is evidence the risk was raised); and the annual record of threat reassessment showing the organization adjusted as information evolved. A documented decision to defer is defensible. Undocumented inaction is not.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

CRYPTO-AGILITY AS END-STATE ARCHITECTURE

Why Crypto-Agility, Not Just PQC

PQC algorithms will eventually need replacement. Cryptographic algorithms have limited lifespans: SHA-1, MD5, DES, 3DES, SSL, and early TLS all fell to attacks or obsolescence. Organizations that hardcoded these algorithms faced expensive emergency migrations. The goal of PQC migration is not merely to swap RSA for ML-KEM; it is to build the organizational capability to change cryptographic algorithms routinely.

Marin's Law on Crypto-Agility: "The effort required to change cryptography in a system is inversely proportional to the agility built into that system from the start." Formally: $Y \approx K / A$, where Y is the migration effort, K is the complexity of the cryptographic estate, and A is the agility level. Systems designed for agility make algorithm changes through configuration updates and rolling deployments. Systems without agility require code rewrites, application rebuilds, and architectural redesigns.

The strongest business case for crypto-agility is that the PQC algorithm set is still evolving. NIST advanced nine additional signature candidates to the third round in May 2026. HQC is being standardized as a non-lattice backup for ML-KEM. Daniel Bernstein's research has raised questions about specific ML-DSA parameter sets. Organizations that hardcode ML-KEM without abstraction will face an emergency migration if lattice-based algorithms are weakened. Organizations that build crypto-agile systems can add HQC as an alternative via configuration. The goal of the PQC migration program is to increase A permanently, so that the next transition takes months rather than years.

Beyond Architecture: Crypto-Agility as an Operational Discipline

The architectural principles below are the easy part. Organizations attempting to implement crypto-agility report that the hard part is the operational, organizational, and process transformation required to make algorithm changes routine.

Crypto-agility is a capability the organization builds across five dimensions. A system with a perfect abstraction layer is not crypto-agile if the organization has no process for deciding when to swap algorithms, no monitoring to verify the swap propagated, and no team trained to evaluate the replacement.

Dimension 1: Architectural agility. Systems are designed so that algorithm selection is configuration-driven, not hardcoded. The test: can the organization change the key exchange algorithm on a Tier-1 service via a configuration change, without modifying

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

application code, rebuilding containers, or filing a change request that takes 90 days to approve?

Dimension 2: Operational agility. The organization has tested, documented processes for executing algorithm changes. Certificate rotation is automated. Monitoring detects drift. Rollback procedures exist and have been tested. The test: has the organization successfully executed an algorithm swap drill, changing the negotiated algorithm for a production service within 10 business days?

Dimension 3: Governance agility. An "approved algorithms" list exists, is maintained, and is enforced through CI/CD checks and network monitoring. The process for evaluating and approving a new algorithm is documented and can be executed within weeks, not months. The test: if NIST published a new algorithm recommendation tomorrow, how long would it take the organization to evaluate, approve, and begin deploying it?

Dimension 4: Skills agility. The organization has people who can evaluate new cryptographic algorithms, assess their suitability, and implement them. This does not require PhD cryptographers on staff. It requires security engineers who understand the provider-pattern architecture, and a CTI function that can interpret cryptanalytic developments. The test: can the organization assess the security implications of a new algorithm within 48 hours of its publication?

Dimension 5: Supply chain agility. Vendors and third parties support configuration-driven algorithm selection. HSM firmware can be updated without hardware replacement. Cloud KMS platforms expose algorithm selection as a customer-configurable parameter. The test: if the organization decided to add HQC support alongside ML-KEM, how many vendor dependencies would block that decision?

Crypto-Agility Architecture Principles

1. **No direct algorithm calls in application code.** Applications use cryptographic providers or adapters that abstract algorithm selection. Examples: Java Cryptography Architecture (JCA) with configurable providers, OpenSSL with pluggable engines, cloud KMS APIs that abstract underlying algorithms.

2. **Algorithm selection is configuration-driven.** A policy file, environment variable, or central configuration service determines which algorithms are used. Changing from X25519 to ML-KEM-768 is a configuration change, not a code change.

3. **Dual-stack capability during transition.** Systems can negotiate both classical and PQC algorithms simultaneously (hybrid mode) and gracefully degrade if the counterparty does not support PQC.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

4. **Automated certificate and key lifecycle management.** Certificate issuance, renewal, rotation, and revocation are fully automated. Manual certificate management at scale is incompatible with the agility requirement.

5. **Algorithm changes are tested in CI/CD.** Cryptographic configuration changes follow the same deployment pipeline as code changes: tested in staging, canary-deployed, monitored, and rollback-capable.

6. **Monitoring of cryptographic posture.** Continuous monitoring detects algorithm drift, expired configurations, and non-compliant deployments. Alerts fire when a system negotiates a deprecated algorithm.

7. **Approved algorithms policy with enforcement.** A maintained list of approved algorithms, enforced through CI/CD policy checks for new code and network monitoring for runtime connections. Unapproved algorithms in Tier-1 production systems trigger remediation.

Crypto-Agility Implementation Roadmap

Crypto-agility is built incrementally, in parallel with the PQC migration. It is not a separate project that starts after migration.

Year 1: Foundation. Mandate crypto-agile design for all new systems (Phase 0 policy). Add crypto-agility requirements to procurement (Phase 7 contract clauses). Assess architectural agility across the estate during discovery (Phase 1), scoring each system for algorithm-change difficulty. Establish the approved algorithms policy and the governance process for updating it. Train the core team on provider-pattern cryptography.

Year 2: Adoption. Implement abstraction layers in Tier-1 services during hybrid migration (Phase 5). Automate certificate lifecycle management (Phase 6). Deploy cryptographic posture monitoring for algorithm drift detection (SOC integration). Conduct the first algorithm swap drill on a non-production system. Update vendor assessment criteria to include crypto-agility support (Phase 7).

Year 3: Verification. Conduct a full production algorithm swap drill: change the negotiated algorithm for a Tier-1 service within 10 business days. Measure and report the Agility KPI. Address systems that failed the drill. Extend crypto-agility monitoring to the full estate.

Year 4+: Maturation. Crypto-agility becomes a standing organizational capability. Algorithm changes are routine operational events. The next cryptographic transition exercises the capability rather than requiring a new multi-year program.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Crypto-Agility OKRs

Objective	Key Result	Measurement Method
Algorithm changes don't require code changes	100% of Tier-1 services use provider-pattern cryptography by Year 3	Architecture review; verified through algorithm swap drill
Certificate rotation is fully automated	Mean time to rotate end-entity certificate < 1 hour; zero manual steps	Automated rotation test across 10 representative services
Algorithm swap executable within 2 weeks	Swap negotiated algorithm for Tier-1 service from hybrid to PQC-only within 10 business days	Annual drill with documented timeline
Cryptographic posture is continuously monitored	100% of Tier-1 endpoints reporting negotiated algorithms to monitoring system	SOC dashboard verification; monthly false-positive review
New algorithm evaluation completed within 48 hours	CTI assessment of new algorithm's relevance and security within 48 hours	Annual simulation exercise
Approved algorithms policy enforced	100% of new deployments pass CI/CD crypto-policy check	CI/CD pass rate; quarterly audit of production cipher suites

Alignment with NIST CSWP 39

In December 2025, NIST published the final Cybersecurity White Paper (CSWP) 39, "Considerations for Achieving Crypto Agility." This publication represents the US government's formal position that crypto-agility is a mandatory organizational capability.

CSWP 39 identifies the same core mechanisms this framework has advocated since its initial release: modular cryptographic architecture with abstraction layers separating algorithms from application logic; configuration-driven algorithm selection; automated certificate and key lifecycle management; continuous monitoring of cryptographic posture; and integration of cryptographic governance into enterprise risk management.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

CSWP 39 also introduces a crypto-agility maturity model, ranging from unstructured and reactive practices to adaptive programs fully integrated into enterprise risk management. This aligns with and validates the framework's maturity model's Crypto-Agility domain assessment.

Where this framework extends beyond CSWP 39 is in three areas. First, Marin's Law provides a quantitative model for measuring crypto-agility that CSWP 39's maturity model does not. Second, the five-dimension model (architectural, operational, governance, skills, supply chain) provides a broader assessment than CSWP 39's primarily architectural focus. Third, the implementation roadmap integrates agility-building into the migration program's existing phase structure rather than treating it as a separate initiative.

NIST's endorsement of crypto-agility as a mandatory capability strengthens every Phase 0 business case that frames PQC migration as a vehicle for building organizational cryptographic agility; the algorithm swap is only the first exercise of that capability.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

REGULATORY & STANDARDS ALIGNMENT MAP

Mapping Framework Phases to Regulatory Requirements

Regulation / Standard	Relevant Phase(s)	What It Requires	Framework Output That Satisfies It
PCI DSS v4.0 Req 12.3.3	Phase 1, 2	Documentation and annual review of cryptographic cipher suites and protocols in use (purpose and location); a full CBOM is the recommended best practice for satisfying this requirement in PQC programs	CBOM; QRA; migration roadmap
NIST IR 8547 (Initial Public Draft)	All	Deprecate quantum-vulnerable public-key algorithms after 2030; disallow after 2035	Full migration execution per roadmap
NIST SP 1800-38 (NCCoE, Preliminary Draft)	Phase 1, 5, 6	Discovery methodology and interoperability/performance testing patterns from the NCCoE Migration to PQC project	Discovery results; pilot reports; compatibility testing
NSA CNSA 2.0	Phase 5, 6, 7	NSS acquisitions CNSA 2.0 compliant by 2027; networking equipment by 2030; web/cloud/OS by 2033; all NSS by 2035	Procurement requirements; vendor compliance verification
EU NIS2 / DORA / CRA	Phase 0, 1, 4	Risk management measures proportionate to risk; supply chain security; incident response	Governance structure; CBOM; vendor governance; evidence dossier
NCSC UK Timelines	Phase 1, 3, 5	Discovery by 2028; high-priority migration by 2031; complete by 2035	Phase 1 complete by 2028; Phase 5 Tier-1/2 by 2031

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ASD (Australia)	All	PQC transition plan by 2026; critical systems migrating by 2028; all Australian Government systems by 2030	Most aggressive timeline — requires immediate action on all phases
OMB M-23-02 (US Federal)	Phase 1, 4	Prioritized inventory of CRQC-vulnerable cryptographic systems with annual updates through 2035; annual funding assessments for migration	Prioritized inventory / CBOM; annual updates; progress reporting

The 2026-2030 Squeeze: Converging Deadlines

Between mid-2026 and end-2030, organizations face at least 14 distinct regulatory deadlines or compliance milestones across multiple jurisdictions. These deadlines compress into a 48-month window.

September 2026: FIPS 140-2 sunset. All FIPS 140-2 certificates move to the Historical list.

November 2026: CMMC Level 2 enforcement begins. FIPS 140-3 validated modules required for CUI protection.

End of 2026: EU Member States should initiate national PQC strategies per the EU PQC Coordinated Roadmap.

January 2027: CNSA 2.0 procurement gate. New NSS acquisitions must include PQC.

March 2027: CA/Browser Forum SC-081v3 reduces maximum certificate lifetimes to 100 days.

2028: NCSC UK initial steps target. ASD (Australia) critical systems migrating.

March 2029: CA/Browser Forum 47-day certificate lifetimes take effect.

2030: CNSA 2.0 software and firmware compliance deadline. EU critical infrastructure transition target. ASD full migration. G7 financial sector critical systems target (2030-2032 window).

2031: Canadian migration target for most systems. NCSC UK high-priority migration deadline.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Organizations that begin their PQC migration program in 2027 will face the mathematical reality that a 4-to-15-year program cannot complete before 2031 at the earliest -- after multiple intermediate deadlines have passed. The window for beginning is now.

New Regulatory Developments Since v1.1 (March 2026)

COM(2026) 13 -- NIS2 PQC Amendment (January 2026): The European Commission proposed explicit PQC transition policy requirements in the NIS2 directive text, closing the interpretation gap where PQC was implied but not mandated. Essential and important entities under NIS2 will need documented PQC transition plans. This strengthens the Phase 0 business case for EU-operating organizations.

G7 PQC Roadmap for Financial Sector (January 2026): Establishes critical financial systems migration by 2030-2032 and full transition by 2035. Direct input to the Financial Services extension roadmap.

NIST CSWP 39 -- Crypto Agility (December 2025, final): Establishes crypto-agility as a mandatory organizational capability. See the Crypto-Agility section for detailed alignment.

NIST Additional Digital Signatures -- Third Round (May 2026): Nine candidates advanced (FAEST, HAWK, MAYO, MQOM, QR-UOV, SDitH, SNOVA, SQIsign, UOV). Evaluation expected to take approximately two years. Not for production planning, but evidence of why crypto-agility is essential.

IETF PLANTS / Merkle Tree Certificates (active 2026): New certificate architecture for post-quantum Web PKI. Google Chrome Quantum-resistant Root Store planned for Q3 2027. Let’s Encrypt committed to MTC issuance June 2026. See Phase 6 PKI Architecture Evolution section.

CA/Browser Forum Ballot SC-081v3: Reduces certificate lifetimes: 200 days (March 2026), 100 days (March 2027), 47 days (March 2029). Drives certificate automation investment that accelerates PQC PKI readiness.

Multinational Regulatory Navigation

Organizations operating across jurisdictions face contradictory hybrid deployment guidance. The recommended approach is a "highest common denominator" strategy:

Deploy hybrid cryptography (classical + PQC) as the default for all Tier-1 traffic. This satisfies jurisdictions that mandate hybrid (BSI Germany, ANSSI France) and is accepted as an interim measure by jurisdictions that prefer pure PQC (NCSC UK, CNSA 2.0).

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Document a pure-PQC transition plan with target dates. This satisfies jurisdictions that treat hybrid as transitional (NCSC UK, CNSA 2.0) and demonstrates intent to all regulators.

When deploying new systems, evaluate whether pure PQC is feasible from the start. NCSC UK's preference for pure PQC on new deployments is technically sound when both endpoints are controlled and the ecosystem supports it.

This approach avoids the operational complexity of maintaining different cryptographic configurations per jurisdiction, while remaining compliant everywhere. Track jurisdiction-specific requirements in the Phase 4 roadmap and adjust the hybrid-to-pure-PQC transition timeline per jurisdiction's expectations.

Algorithm Sovereignty and Standards Fragmentation

The hybrid-mandate divergence above is the visible part of a larger fragmentation: jurisdictions are diverging on which PQC algorithm families they will accept at all. South Korea has selected its own KpqC algorithms alongside the NIST suite; China is expected to extend its SM national-algorithm series with domestic PQC selections rather than adopt NIST's; Russia maintains the GOST track; and European sovereignty initiatives press for European-controlled implementations. For a multinational, this is no longer a hybrid-posture question. It is an algorithm portfolio question.

Three planning consequences follow. Crypto-agility requirements harden: the architecture must support multiple PQC families concurrently, selected per jurisdiction by policy (precisely the provider-pattern abstraction the Crypto-Agility foundation specifies), and single-family assumptions in new systems should be treated as a design defect. The vendor questionnaire gains a question: which non-NIST national algorithm families (KpqC, SM-series, GOST) the product supports or has on its roadmap, for organizations operating in those jurisdictions. And the Phase 4 roadmap gains jurisdiction-specific algorithm tracks: the same logical migration may require different algorithm deployments per region, with the regulatory intelligence function (GRC Implementation) tracking each jurisdiction's selections as they finalize.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

SKILLS & TEAM STRUCTURE

The Skills Challenge

PQC migration requires a combination of skills that rarely coexists in a single team: deep cryptographic knowledge (algorithms, protocols, PKI architecture), enterprise program management at scale (governance, stakeholder management, multi-year planning), and domain-specific technical expertise (network security, application security, OT engineering, cloud architecture). The market for professionals who combine even two of these is thin.

The realistic approach is not to hire a complete team of PQC specialists. It is to build the program around a small core of cryptographic expertise, supplemented by existing enterprise security and IT staff who are upskilled on PQC-relevant topics, and augmented by external specialists for capabilities the organization cannot develop internally in the required timeframe.

Core Roles

Role	Skills Required	Source	Typical FTE
QRPM	Program management; stakeholder management; risk governance; basic crypto literacy; board-level communication	Internal senior PM with PQC training	1.0
Cryptographic Architect	Deep cryptography expertise; PKI architecture; protocol design; PQC algorithm knowledge; crypto-agility design patterns	Internal security architect + specialized training; rare external hire	0.5–1.0
Security Engineers (PQC)	TLS/SSH/IPsec configuration; HSM management; certificate lifecycle; library evaluation; hybrid deployment	Internal security engineers with PQC training	2–4

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Application Security Lead	Code review; SAST/DAST tooling; library management; CI/CD pipeline integration; crypto-agility patterns	Internal AppSec team with crypto-agility focus	1.0
OT Security Specialist	ICS/SCADA knowledge; OT network architecture; safety case management; vendor coordination	Specialized hire or external partner; very scarce	0.5–1.0 (if OT in scope)
Vendor/Procurement Lead	Contract negotiation; RFP management; vendor relationship management; SLA design	Internal procurement with PQC requirements training	0.5
PMO Analyst	KPI tracking; reporting; evidence dossier management; SteerCo coordination	Internal PMO or shared resource	0.5–1.0

Team Sizing

Team size depends on the cryptographic estate's complexity more than on organizational revenue or headcount. A 50,000-employee financial institution with 2,000 TLS endpoints, 15 HSMs, and 200 vendor relationships requires a larger program team than a 200,000-employee manufacturer with 500 TLS endpoints concentrated in a single ERP platform.

Sizing heuristic: one dedicated FTE per 500 cryptographic instances in the CBOM for the first two years (discovery, CBOM, risk scoring, and pilot phases). This drops to one per 1,000 during production rollout as processes mature and tooling automates repetitive tasks. The QRPM, Cryptographic Architect, and PMO Analyst are dedicated overhead regardless of estate size.

For organizations with fewer than 1,000 cryptographic instances, a part-time QRPM with consulting augmentation for the Cryptographic Architect role is viable. For organizations exceeding 10,000 instances, plan for a dedicated program office with 8–12 FTEs at peak.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Skills Matrix

The skill domains below map to the work the phases generate. Use the matrix to assign named individuals per domain and to target upskilling; a domain with no name against it is a program risk.

Skill domain	Primary roles	Proficiency target
Program governance and strategy	QRPM, executive sponsor, SteerCo members	Can defend the roadmap, budget, and risk acceptances to the board and regulators
Cryptographic discovery and CBOM	Security architects, Inventory & Discovery workstream lead	Can operate multi-layer discovery tooling and maintain a queryable CycloneDX CBOM
PKI, KMS, and HSM engineering	PKI engineers, infrastructure security	Can execute dual-stack CA operation, HSM PQC migration, and key ceremonies
Protocol and application engineering	Application security leads, platform engineers	Can implement and debug hybrid TLS/IPsec/SSH and cryptographic library migrations
Assurance, testing, and performance	Security testing and performance engineers	Can design negative tests, interoperability matrices, and PQC performance baselines
OT and embedded	OT security, embedded engineers	Can apply gateway and compensating-control patterns within safety constraints

Detailed curricula, candidate roles, and evaluation criteria for each domain are maintained in the source article library (see Accompanying Resources).

Build, Borrow, or Buy

Capability	Recommended Model	Rationale
Program management	**Build.** Internal QRPM.	Institutional knowledge and continuity matter for multi-year programs. Acceptable to borrow for

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

		the first 6 months while identifying an internal QRPM.
Cryptographic architecture	**Build if possible; borrow for design.**	This capability has lasting value beyond PQC. Use consultants to design the architecture; train internal staff to maintain it.
Discovery and inventory	**Buy the tool; run it internally.**	Internal staff understand the estate better than any external team. Acceptable to borrow for initial deployment and configuration.
PKI modernization	**Build, augmented as needed.**	PKI errors cause outages. Production CA operations require people who understand the production environment. Borrow for architecture design and migration planning.
Vendor governance	**Build.** Internal procurement.	Vendor relationships are organizational assets. Acceptable to borrow for developing questionnaire frameworks and assessment criteria.
Strategic quantum CTI	**Borrow** for most organizations.	The skill set (interpreting resource estimation papers, tracking national quantum programs) is specialized and scarce. Contract quarterly strategic assessments from a qualified provider.

Training Approach

Training operates at four levels:

1. Executive education (half-day to one day). Audience: SteerCo members, board risk committee, senior leadership. Content: quantum threat in business terms, regulatory deadlines, program governance responsibilities, KPI interpretation. Outcome: informed sponsors who can approve risk appetite statements and budget commitments.

2. PQC foundations (3–5 days). Audience: all workstream participants, security engineers, architects, application developers with cryptographic touchpoints. Content: algorithm overview (ML-KEM, ML-DSA, SLH-DSA, FN-DSA), hybrid deployment mechanics, CBOM concepts, risk assessment methodology, crypto-agility design patterns. Outcome: a team that can execute Phase 1–3 activities without constant expert supervision.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

3. Deep technical (ongoing, lab-based). Audience: security engineers and architects who will configure, test, and deploy PQC. Content: hands-on exercises with hybrid TLS deployment, HSM PQC configuration, CBOM generation using CycloneDX, certificate lifecycle automation, performance testing methodology. Outcome: practitioners who can run pilots and production deployments.

4. Crypto champion program. Designate one "crypto champion" per platform or application team (web, mobile, data, infrastructure, OT, identity). Champions attend PQC foundations training, participate in quarterly crypto-agility briefings, and serve as the liaison between the PQC program and their platform team. Champions sign off on "crypto readiness" in design reviews for new systems and shepherd PQC library upgrades within their domain. The champion model scales the program's reach without requiring every developer to become a cryptography specialist.

Curricula, certification paths, and self-study sequences for each training level are covered in the companion book, *Quantum Ready* (QuantumReady.com).

Sustaining Capability Beyond the Migration

The skills and structures built for PQC migration must outlast the program itself. Plan for the program to transition from a dedicated initiative to a standing organizational capability:

- The QRPM role evolves into a permanent "Cryptographic Governance Lead" within the CISO's function
- Crypto champions become a standing network, analogous to security champion programs
- The cryptographic inventory and CBOM become maintained operational assets
- Crypto-agility OKRs become standing metrics in the enterprise security dashboard
- SOC detection rules for cryptographic posture monitoring become part of the permanent detection library

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

SOC IMPLEMENTATION

PQC migration creates operational security responsibilities that extend beyond the migration program. The Security Operations Center must develop detection capabilities for quantum-related threats, build a threat intelligence function that tracks CRQC developments and PQC implementation vulnerabilities, and maintain incident response playbooks for scenarios that have no equivalent in the SOC's current library. These capabilities are needed during the migration (to verify that migrated systems stay migrated and that hybrid implementations are not silently downgraded) and permanently after migration is complete (to maintain cryptographic posture, detect drift, and respond to the inevitable future algorithm transitions).

This section specifies what the SOC must build, in what order, and with what resources. It is designed for SOC directors and senior analysts who are familiar with enterprise detection engineering and need to understand how quantum security maps to their existing infrastructure.

Prerequisites

Every SOC detection capability described below depends on a single prerequisite: the SOC must have access to a queryable, continuously updated cryptographic posture registry that maps systems to their expected cryptographic configuration. For each system in scope, the registry must specify: whether the system should be running hybrid, classical-only (pending migration), or PQC-only; which specific algorithms and cipher suites are expected; and when the system's migration status last changed.

This registry is a derivative of the Phase 1/Phase 2 cryptographic inventory and CBOM. The migration program builds it. The GRC function governs it. The SOC consumes it. If the registry lives in a GRC spreadsheet updated quarterly and shared by email, the detection rules described below cannot function. The registry must be machine-readable and integrated with the SIEM, ideally through an API or automated feed that updates as migration progresses. Designing this integration should be a Phase 1 architecture decision.

Detection Use Cases

The SOC's contribution to quantum security centers on five detection capabilities that can be built on existing infrastructure: the SIEM, network monitoring tools, and certificate management systems already in place. SOC engineers familiar with TLS monitoring and

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

DLP rules will recognize these as extensions of what they already do. None requires exotic quantum-specific technology, though each requires updates to protocol parsing and correlation rules.

Use Case 1: Hybrid Downgrade Detection

Organizations deploying hybrid PQC implementations face a specific attack vector: an adversary forcing a connection to fall back to classical-only cryptography. This is analogous to the TLS downgrade attacks SOCs have monitored for years.

Detection approach. SIEM correlation rules that flag connections negotiating classical-only key shares (named groups) when the expected configuration is hybrid. A TLS 1.3 session between two systems that both support ML-KEM should not complete a handshake using only ECDH. If it does, that is either a misconfiguration or an active downgrade attack.

Illustrative detection rule. Alert when a TLS handshake between two endpoints listed as "hybrid-enabled" in the cryptographic posture registry completes using a key share that does not include an ML-KEM or other PQC key exchange component. Severity: High. Expected false positive rate in a well-maintained registry: 5–10%. In a poorly maintained registry: high, because stale entries generate constant noise.

Protocol parsing requirement. In TLS 1.3, hybrid key exchanges are negotiated using IANA-assigned NamedGroup codepoints, not X.509 OIDs. The NamedGroup values for ML-KEM-768, ML-KEM-1024, and the hybrid key exchange groups must be added to the SOC's traffic analysis rules manually. As of mid-2026, this is custom engineering work in most SIEM platforms. Plan for 1–2 weeks of rule development and testing per monitoring platform.

Middlebox-induced false positives. The most common source of false positives will be the organization's own infrastructure. PQC key encapsulations are larger than classical ECDH ephemeral keys, and hybrid TLS ClientHello messages frequently exceed the size that legacy firewalls, secure web gateways, and SSL inspection proxies can handle without fragmentation. These middleboxes may drop fragmented handshakes or strip unrecognized extensions, forcing a classical-only fallback that looks identical to a downgrade attack. Detection rules must account for benign middlebox-induced downgrades, and the migration program (Phase 6) should flag middlebox upgrades as a prerequisite for reliable hybrid deployment.

Use Case 2: Cryptographic Drift Monitoring

Migration is not a one-time event. After an organization migrates systems to PQC, those systems drift. A developer deploys a new microservice using a classical-only TLS library because that is what the tutorial showed. An IT team restores a database server from a

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

pre-migration backup. A SaaS integration uses an API endpoint that silently reverts to RSA key exchange after a vendor update. Shadow IT operates entirely outside the migration program's scope.

Detection approach. Continuous network monitoring that flags connections using deprecated cipher suites that the migration program has marked for retirement. This is continuous compliance verification at the network layer, operating in real time rather than in quarterly audit cycles.

Illustrative detection rule. Alert when any internal endpoint initiates or accepts a TLS session using a classical-only key share with a system classified as "migration complete" in the posture registry. Severity: Medium (likely misconfiguration), escalating to High if the same endpoint generates repeated classical-only sessions across multiple connections.

Illustrative metric. Percentage of monitored TLS sessions using PQC or hybrid cipher suites versus classical-only, measured weekly. Tracked as "Cryptographic Migration Coverage" on the SOC dashboard. The trend should be monotonically increasing toward 100% for in-scope systems, with any regression investigated within 24 hours.

Visibility requirement. Measuring this accurately requires visibility into internal east-west traffic, not just north-south perimeter flows. Many SOCs have limited internal traffic inspection capability, particularly in cloud-native environments where service mesh encryption happens at the application layer. Extending visibility to internal traffic may require infrastructure investment beyond the SOC's budget authority. This cost should be included in the Phase 0 business case.

Use Case 3: Certificate Lifecycle Anomalies

PQC migration involves wholesale transition of certificate infrastructure: new certificates using PQC signature algorithms (ML-DSA, SLH-DSA), reconfigured certificate chains, and new intermediate Certificate Authorities supporting the PQC hierarchy alongside the legacy one. This transition creates a window of elevated risk.

Detection targets. Failed certificate chain validations involving PQC certificates, which may indicate interoperability issues or deliberate manipulation. Severity: Medium, with investigation required to determine root cause. Unauthorized certificate issuance during the transition period. Alert on any certificate issuance event from the organization's internal CA that uses a PQC signature algorithm not on the approved list. Severity: High. CA signing key access outside scheduled certificate issuance windows, particularly during the transition period. Severity: Critical.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Tooling gap. Certificate transparency monitoring for PQC certificates is still immature. Most CT log monitors are built for the classical certificate ecosystem. The SOC may need to build custom monitoring around the organization's internal CA logs rather than relying on external CT infrastructure.

Use Case 4: TNFL and Signature Integrity Monitoring

The quantum threat to digital signatures (Trust Now, Forge Later) enables an adversary to take actions in the present: forging software updates, fabricating financial instructions, impersonating trusted parties. Unlike HNDL, which compromises data created in the past, TNFL enables real-time exploitation once a quantum adversary has signature forgery capability.

Detection approach. Heightened monitoring of code signing events, software update authentication, and systems that depend on automated signature verification for trust decisions (firmware updates, financial transaction authorization, identity federation).

Illustrative detection rules. Alert on code signing events using unexpected signing keys, unexpected signing identities, or signing events at unusual times. Severity: Critical. Alert on any modification to signature verification policies in production systems. Severity: Critical. Monitor for anomalous signing volume: a sudden spike in certificate signing requests or code signatures may indicate automated exploitation of a compromised signing key. Establish a baseline and alert on deviations exceeding two standard deviations. Severity: High.

Illustrative metric. Mean time to detect an unauthorized signing event (MTD-Signing). Target: under 15 minutes for code signing events associated with production deployments. Measured through periodic red team exercises.

Organizational challenge. Most organizations do not have centralized visibility into all signing events. Code signing may be managed by the development team, firmware signing by the OT team, and certificate signing by the PKI team, each with separate logging and alerting. Unifying these into a single SOC-monitored view requires cross-functional cooperation.

Use Case 5: Enhanced HNDL-Indicator Detection

The HNDL threat is active today. Adversaries are intercepting and storing encrypted data for future quantum decryption. The SOC's role is to detect indicators of harvesting activity weighted by the sensitivity and longevity of the data being targeted.

Detection approach. This is the same data exfiltration monitoring SOCs already perform, but with quantum-informed reprioritization. A slow, sustained exfiltration of archived diplomatic correspondence or pharmaceutical R&D data is potentially more damaging in

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

a quantum context than a one-time dump of transactional data, because long-lived data retains its value across the decryption timeline.

Illustrative detection rules. Alert on sustained outbound data transfers from network segments hosting data classified as having secrecy requirements exceeding 10 years, even if the transfer volume per session is below traditional DLP thresholds. Severity: High. Alert on bulk access patterns to archival storage by service accounts or user accounts that do not normally access those systems. Severity: High.

Illustrative metric. "HNDL Exposure Score": a composite metric tracking the percentage of high-secrecy-requirement data stores with active DLP coverage, the mean time to detect anomalous access to those stores, and the percentage of data classified with secrecy-requirement metadata. Target: 100% coverage for data stores with 10+ year secrecy requirements, with mean detection time under 4 hours.

Dependency. HNDL detection depends entirely on data classification quality. If the organization has not classified its data by secrecy requirement duration, the SOC cannot weight alerts by HNDL relevance. The classification exercise belongs to the information governance function, not the SOC, but the SOC should flag the gap if it exists and escalate it through the KRI framework.

Cyber Threat Intelligence

Large enterprise SOCs with integrated CTI teams operate across three horizons. Quantum threat intelligence fits naturally across all three, but each level requires different sources, analytical skills, and outputs.

Tactical CTI (Hours to Days)

PQC implementation vulnerabilities. When a CVE is published against a specific version of liboqs, or a side-channel vulnerability is discovered in a hardware vendor's ML-KEM implementation, the CTI team assesses exposure and feeds the assessment to the SOC for immediate triage.

Cryptanalysis papers with implementation impact. When researchers identify a weakness in a specific parameter set, the CTI team evaluates whether the organization's deployed implementations are affected and produces an actionable assessment within hours.

Illustrative metric. Time from PQC vulnerability disclosure to CTI assessment delivered to SOC (TTAssess-PQC). Target: under 4 hours for vulnerabilities affecting deployed

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

implementations; under 24 hours for vulnerabilities in algorithms the organization plans to adopt.

Required sources. NIST PQC mailing list, vendor security advisories, IETF working group notifications (particularly pquip and lamps groups), CERT advisories, and curated cryptanalysis feeds.

Operational CTI (Weeks to Months)

Changes in adversary collection patterns consistent with HNDL acceleration. If a tracked APT group shifts targeting from operational intelligence (short-lived value) toward strategic archives and R&D repositories (long-lived value), that behavioral change is potentially significant in a quantum context.

Supply chain indicators relevant to PQC. Anomalous contributions to open-source PQC implementations. Probing of PQC library supply chains.

Honest assessment. The operational CTI use cases for quantum are speculative at this stage. No publicly attributed campaign has been documented as specifically targeting PQC implementations for supply chain compromise (as of mid-2026). This is a watch-and-wait posture, not an active hunt.

Strategic CTI (Months to Years)

Tracking quantum hardware progress toward CRQC. Monitor developments across quantum computing modalities (superconducting, trapped ion, photonic, neutral atom, silicon spin) and assess implications against CRQC Quantum Capability Framework dimensions: error correction, logical gate operations, decoder performance, continuous operation, and engineering scale.

Monitoring national quantum programs. State investment in quantum computing is the primary driver of CRQC development. Track the progress of programs in the US, China, the EU, Australia, Japan, South Korea, and others relevant to the organization's threat model.

Monitoring cryptanalytic research trends. Resource estimation papers and novel algorithmic approaches to cryptanalysis have direct implications for migration urgency.

Output. A quarterly "Quantum Threat Landscape Assessment" delivered to the CISO and the risk committee. The assessment covers: changes to CRQC timeline estimates, regulatory deadline updates, cryptanalytic developments affecting deployed or planned algorithms, and recommendations for adjusting migration priorities. This assessment feeds the Material Developments KRI at the board level.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Skills challenge. Producing useful strategic quantum CTI requires reading arXiv preprints, understanding quantum error correction papers, interpreting resource estimation models, and evaluating hardware announcements against a technically grounded framework. Three realistic options: hire a specialist (expensive and scarce), train an existing analyst (6–12 months), or contract the analysis externally (pragmatic for most organizations in the near term). The Build, Borrow, or Buy guidance in the Skills & Team Structure section recommends borrowing this capability for most organizations.

Incident Response Playbooks

Four quantum-specific playbooks should be developed during Phase 0 governance setup. Each specifies a trigger, immediate actions, coordination requirements, and a drill metric.

Playbook 1: PQC Algorithm Vulnerability Disclosure

Trigger. NIST, a CERT, or a credible research group publishes a vulnerability or weakness in a PQC algorithm or implementation deployed in the organization's environment.

Immediate actions (0–4 hours). CTI team assesses severity and scope. Is this a full algorithm break (catastrophic) or an implementation bug in a specific library version (serious but contained)? SOC queries the cryptographic posture registry to identify all systems running the affected algorithm or library. Initial impact assessment delivered to the CISO.

Short-term actions (4–48 hours). If implementation-specific: initiate emergency patching, prioritizing internet-facing and high-value systems. If the vulnerability affects the algorithm itself: activate the organization's crypto-agility capabilities to begin algorithm rotation.

Coordination. The SOC cannot execute this playbook alone. Pre-established coordination between the SOC, the PQC migration program office, the application teams, and the CISO. The playbook should specify a named incident commander, a communication chain, and pre-authorized emergency change windows.

Drill metric. Time from vulnerability disclosure to completed impact assessment (all affected systems identified). Target: under 4 hours.

Playbook 2: Confirmed Hybrid Downgrade Attack

Trigger. SOC detects and confirms (not a false positive) that a connection between two PQC-capable systems has been downgraded to classical-only cryptography by an external adversary.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Immediate actions. Isolate the affected network path. Preserve packet captures for forensic analysis. Determine the mechanism: a TLS inspection proxy or middlebox stripping PQC extensions, a compromised endpoint advertising only classical key shares, or an active attacker exploiting a fallback vulnerability. Assess scope: are other connections being downgraded simultaneously?

Escalation. A confirmed downgrade attack is a strong indicator of an active, sophisticated adversary targeting the organization's cryptographic infrastructure. Escalate to the CISO immediately.

Playbook 3: Credible CRQC Announcement

Trigger. A credible entity announces the achievement of a CRQC capable of running Shor's algorithm against production cryptographic key sizes.

Immediate actions. Invoke crisis communication procedures. Produce an impact assessment: what percentage of systems have been migrated? What data protected by classical-only algorithms has the highest sensitivity? What systems dependent on classical digital signatures are most vulnerable to TNFL exploitation? Activate emergency migration procedures for the highest-priority remaining systems.

Context. This is the playbook where prior preparation matters most and in-the-moment response matters least. The right time to have responded to this scenario was years earlier, through completing the migration program.

Playbook 4: Emergency Algorithm Rotation

Trigger. The organization initiates an emergency rotation of cryptographic algorithms.

SOC's role. The SOC's primary responsibility during an algorithm rotation is to not mistake the rotation for an attack. An authorized emergency rotation produces mass certificate revocations, sudden cipher suite changes, and key management system activity spikes. The SOC must be briefed on the rotation schedule, with suppression rules for expected activity, while maintaining detection for actual adversary activity.

Drill metric. False positive rate during algorithm rotation drills. Target: less than 10%.

Tabletop Exercises

Playbooks exist on paper until they are exercised. The following exercises should be run at least annually, with the PQC migration program office, the CISO, GRC function, and relevant application owners participating.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Exercise 1: The Friday Afternoon CVE. A critical CVE is published at 16:30 on a Friday against the specific version of liboqs deployed in the organization's production TLS termination layer. Tests: Can the SOC identify all affected systems within 4 hours? Who authorizes emergency patching outside the normal change window?

Exercise 2: The Ambiguous Cryptanalysis Paper. An arXiv preprint from a credible research group claims a polynomial-time attack against ML-KEM at security level 3. Within 24 hours, cryptographers disagree on the claim's validity. Tests: How does the CTI team evaluate credibility? At what point does the SOC escalate?

Exercise 3: Silent Downgrade. A cluster of internal services have been negotiating classical-only TLS for 72 hours, coinciding with a vendor load balancer firmware update. Tests: Was the downgrade detected automatically? Was it a misconfiguration or adversary exploitation?

Exercise 4: The CRQC Announcement. A credible national laboratory announces Shor's algorithm factoring a 1024-bit RSA key in under 8 hours. Tests: communication chain, exposure assessment speed, emergency migration readiness, pre-drafted external communications.

Exercise 5: Signing Key Compromise. The SOC detects an unauthorized code signing event; the signed firmware has already been distributed to three branch offices. Tests: artifact identification, rollback procedures, key revocation without disrupting legitimate signed artifacts.

For each exercise, produce a structured after-action report: what worked, what broke, what assumptions proved incorrect, and what changes are required.

SOC Metrics Summary

Metric	What It Measures	Target
Cryptographic Migration Coverage	% of monitored TLS sessions using PQC or hybrid	Monotonically increasing; regression investigated within 24 hours
Cryptographic Drift Incidents	Systems reverting to classical-only after migration	Zero; any non-zero value triggers investigation

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

TTAssess-PQC	Time from PQC vulnerability disclosure to completed CTI assessment	Under 4 hours for deployed implementations
Impact Assessment Speed	Time from disclosure to identification of all affected systems	Under 4 hours
MTD-Signing	Mean time to detect unauthorized signing events	Under 15 minutes for production signing keys
HNDL Exposure Score	DLP coverage + detection latency + data classification completeness	100% coverage for 10+ year secrecy data stores
Algorithm Rotation Drill FP Rate	False positive rate during rotation exercises	Under 10%

Skills and Tooling Gaps

The protocol parsing gap. Most SIEM platforms, network detection tools, and TLS inspection devices cannot yet parse PQC cipher suite identifiers, hybrid key exchange groups, or PQC-specific TLS extensions. Plan for 2–4 weeks of engineering effort to build and test custom Suricata rules, Zeek scripts, or SIEM enrichment pipelines, with ongoing maintenance as PQC implementations evolve.

The CTI skills gap. Tactical PQC CTI (monitoring advisories, tracking library updates) can be handled by existing analysts with modest training. Strategic CTI (evaluating quantum hardware milestones, interpreting resource estimation papers) requires specialized knowledge. The realistic path for most organizations: contract strategic quantum CTI externally and build tactical and operational capability in-house over 12–18 months.

The east-west visibility gap. Many SOCs have limited internal traffic inspection capability, particularly in cloud-native environments. Extending east-west visibility is a prerequisite for cryptographic drift monitoring and may require infrastructure investment.

SOC Implementation Roadmap

Phase 1 (0–3 months). Operationalize the cryptographic posture registry for SOC access. Stand up tactical CTI monitoring of PQC vulnerability advisories. Draft the PQC Algorithm Vulnerability Disclosure playbook and conduct the first tabletop exercise.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Phase 2 (3–9 months). Implement hybrid downgrade detection rules for the highest-priority network segments. Build or acquire PQC cipher suite parsing capability. Begin cryptographic drift monitoring for migrated systems. Conduct the first HNDL-focused review of data classification and DLP coverage.

Phase 3 (9–18 months). Extend detection rules to full enterprise scope. Integrate certificate lifecycle monitoring. Implement TNFL-specific signing event monitoring. Establish the quarterly Quantum Threat Landscape Assessment. Conduct a full algorithm rotation drill.

Phase 4 (Ongoing). Refine detection rules based on operational experience. Update playbooks based on exercise findings. Track and report SOC quantum security metrics as part of regular SOC reporting. Adapt as PQC implementations, standards, and the threat landscape evolve.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

GRC IMPLEMENTATION

The Governance, Risk, and Compliance function owns the risk management infrastructure that makes PQC migration governable, measurable, and auditable. Where the SOC detects and responds, the GRC function measures, reports, and governs. The GRC function produces the instruments that the board uses to oversee the program, the CISO uses to manage it, and internal audit uses to verify it: the risk appetite statement, the KRI framework, the regulatory intelligence process, the vendor quantum readiness assessments, and the evidence dossier.

This section specifies what the GRC function must build, the outputs it must produce, and how those outputs connect to the SOC's detection capabilities and the migration program's execution.

Why Quantum Risk Breaks the Standard ERM Playbook

Enterprise risk management frameworks are designed around risks that have observable precursors, historical frequency data, established loss models, and reasonably bounded timelines. Quantum risk has none of these.

The timing is uncertain. Estimates for a cryptanalytically relevant quantum computer range from 10 to 30+ years, with probability distributions wide enough to paralyze a risk committee accustomed to tighter confidence intervals. There are no historical incidents to calibrate against. The risk straddles categories: it is simultaneously a technology risk (can we migrate in time?), a compliance risk (are we meeting regulatory deadlines?), a strategic risk (are our competitors getting ahead?), and an operational risk (will the migration disrupt business?).

The standard ERM response (log the risk with a "Low likelihood / High impact" rating, assign it to the CISO, review annually) is risk acknowledgment, not risk management. There is a large gap between the two.

The fix is not a separate quantum risk framework. It is adapting the existing ERM framework so that quantum risk receives the same structured treatment (risk appetite statements, KRIs, escalation thresholds, board reporting cadence) that the organization applies to its other top-tier risks.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Risk Appetite and Tolerance

A quantum risk appetite statement translates the board's intent into decision criteria the program can act on. Without it, every trade-off between migrating a complex legacy system (expensive, disruptive) and deferring it (cheap, smooth, but leaves exposure) must be escalated because there is no governance instrument defining which choice aligns with the board's intent.

Drafting the Statement

The statement operates at two levels.

Strategic level (board language). A single sentence that defines the organization's overall posture. Example: "The organization will complete migration of all systems protecting data with confidentiality requirements exceeding 10 years before the earliest credible CRQC estimates, and will achieve compliance with all applicable PQC regulatory deadlines with a minimum 12-month buffer."

Operational tolerances (program language). Specific, measurable thresholds for each risk dimension:

HNDL exposure tolerance. "No more than 20% of data classified as having secrecy requirements exceeding 10 years will remain protected by quantum-vulnerable algorithms by end of 2027. The target is 0% by end of 2029."

TNFL exposure tolerance. "All production software and firmware signing will use NIST-approved quantum-resistant signatures (SP 800-208 hash-based or ML-DSA, dual-signed with classical) by end of 2027. All certificate authority signing keys will be PQC-capable, with hybrid or PQC certificate issuance available, by end of 2029."

Regulatory compliance tolerance. "The organization will achieve compliance with all PQC-related regulatory requirements at least 12 months before the applicable deadline. Zero tolerance for deadline non-compliance."

Crypto-agility tolerance. "All newly deployed systems from Q3 2026 onward must implement crypto-agile architecture as a mandatory design requirement. No exceptions without Steering Committee approval and a documented remediation plan."

These thresholds are illustrative. The right values depend on the organization's sector, regulatory environment, data profile, and risk culture. The requirement is that they exist, are specific enough to drive decisions, and are reviewed annually.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Approval and Review

The risk appetite statement should be presented to the board for approval during Phase 0, alongside the program charter and budget request. It should be reviewed annually at the board level, incorporating: changes to CRQC timeline estimates (from the quarterly Quantum Threat Landscape Assessment produced by the CTI function, as described in the SOC Implementation section), new regulatory deadlines, and migration progress.

Key Risk Indicators: Three-Level Cascade

KRIs make governance operational. For PQC migration, they cascade across three organizational levels. Each level serves a different audience at a different reporting frequency.

Board-Level KRIs (Quarterly)

The board needs four or five indicators that answer the questions directors actually ask: Are we on track? Are we compliant? What is our exposure? What has changed?

KRI	What It Answers	Illustrative Threshold
Migration Completion Rate	Are we on track?	Variance >5 percentage points from plan triggers SteerCo review and board notification
HNDL Residual Exposure	How much sensitive data is harvestable now?	Any quarter where this fails to decrease triggers risk committee escalation
Regulatory Compliance Buffer	Do we have margin?	Buffer <12 months: amber; <6 months: red and board intervention
Third-Party Quantum Readiness	Are our vendors keeping up?	Any Tier 1 vendor below "planning" triggers formal engagement
Material Developments Flag	Has the risk picture changed?	Qualitative: produced by CTI function, packaged by GRC for board

CISO-Level KRIs (Monthly)

KRI	What It Reveals	Illustrative Threshold

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Inventory Completeness	Is our migration plan based on a complete map?	Below 80%: migration plan is unreliable; priority shifts to inventorying
Migration Velocity	Have we hit a blocker?	Velocity <70% of plan for two consecutive months triggers program review
Crypto-Agility Adoption	Are we creating new technical debt?	Any month below 95% triggers architecture governance review
Regulatory Horizon Tracker	What is coming?	Count of pending requirements with compliance status for each
Vendor Readiness Scores	Are vendors regressing?	Any Tier 1 vendor dropping a level triggers engagement review
Budget Consumption vs. Plan	Is the program executing or stalling?	Underspend >15%: triggers review (usually signals blocked workstreams)

Operational KRIs (Weekly / Real-Time)

KRI	What It Catches	Illustrative Threshold
Weekly Migration Throughput	Bottlenecks before they compound	Instances migrated in past 7 days, by system tier
Cryptographic Drift Count	Migrated systems reverting	Any non-zero value investigated; persistent drift = systemic issue
Hybrid Downgrade Alerts	Active attacks or misconfiguration	Broken down by confirmed attacks, misconfigs, and false positives
Certificate Transition Progress	PKI migration tracking	% of certificates reissued with PQC signature algorithms
Open Exceptions/Deferrals	Governance drift	Any exception >6 months without remediation plan triggers escalation

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

TTR-PQC	Remediation speed	Time from PQC vulnerability disclosure to confirmed remediation

Regulatory Intelligence

PQC regulation is moving faster than most GRC teams appreciate. A GRC function that reviews regulatory developments annually is reviewing it too infrequently. Quarterly is the minimum cadence. For organizations in regulated sectors, monthly monitoring is more appropriate.

Building the Process

The regulatory intelligence function does not require a dedicated team. It requires a defined process, assigned ownership, and a structured output.

Sources to monitor. NIST PQC publications and the NCCoE migration project. NSA CNSA 2.0 updates. EU NIS Cooperation Group outputs and the coordinated PQC implementation roadmap. CISA advisories and product category updates. Sector-specific regulators (BIS/BCBS for banking, EMA for pharma, NERC for energy). Standards bodies (ETSI, ISO, IETF). National cybersecurity agencies in jurisdictions where the organization operates.

Output. A quarterly Regulatory Horizon Report delivered to the Steering Committee and the risk committee. The report lists each regulatory development, its jurisdiction, effective or expected date, implications for the migration plan, and required action. Developments are classified as: Confirmed (enacted, final rule), Proposed (in legislative or regulatory process), or Signaled (guidance, recommendations, speeches by senior officials indicating direction).

The report should flag divergences between jurisdictions that affect the organization.

Escalation criteria. A new confirmed requirement that affects the organization's migration timeline triggers an immediate Steering Committee assessment, outside the quarterly cycle. A proposed requirement triggers an impact assessment within 30 days. A signal triggers a monitoring entry for the next quarterly review.

Third-Party Quantum Readiness Assessment

An organization can migrate its own cryptographic infrastructure flawlessly and still be exposed if its critical vendors, partners, and service providers have not. The GRC function operationalizes this risk on an ongoing basis, complementing the Phase 7 vendor governance process.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Vendor Tiering

Tier 1 (migration-constraining). Vendors whose PQC readiness directly constrains the organization's migration timeline: HSM vendors, certificate authorities, cloud infrastructure platforms, payment processors, managed security services providers. Require detailed assessment and active engagement.

Tier 2 (cryptographic dependency). Vendors whose products or services involve cryptographic operations but do not directly constrain migration timing: SaaS platforms handling sensitive data, API partners, data analytics platforms with encryption at rest. Require standard assessment.

Tier 3 (no cryptographic dependency). Vendors with no cryptographic touchpoint. No PQC assessment required.

Assessment Mechanics

Incorporate quantum readiness into the existing third-party risk assessment process. The assessment should cover: whether the vendor has conducted a cryptographic inventory of the products or services they provide; whether they have a PQC migration roadmap with dates; whether they support hybrid or PQC cipher suites in current product versions; and whether their own subcontractors are also preparing.

For Tier 1 vendors, require a Cryptographic Bill of Materials (CBOM) in a machine-readable format such as CycloneDX 1.6+, which added native CBOM support. If a vendor cannot produce a CBOM, that itself indicates a level of cryptographic visibility that should concern the organization.

For Tier 1 vendors, the assessment should also include a meeting with the vendor's product security team to review their PQC roadmap and validate their algorithm support claims.

Contractual Provisions

New and renewed contracts with cryptographic dependencies should include: a requirement to support NIST-standardized PQC algorithms by a specified date; notification obligations for any cryptographic vulnerability affecting the contracted product or service; a right to audit the vendor's cryptographic practices; and a commitment to support crypto-agile configurations that allow algorithm changes without contract renegotiation.

Concentration Risk

If the organization's PQC migration depends on a single HSM vendor, a single certificate authority, or a single cryptographic library, that concentration is a risk. For each Tier 1 vendor with a cryptographic dependency, maintain a documented substitution

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

assessment: if this vendor cannot deliver PQC support by the required date, what is the alternative? How long would a substitution take? What is the cost?

M&A Due Diligence

Mergers and acquisitions are the special case of third-party risk where the third party becomes first-party: the acquirer inherits the target's cryptographic estate and its unremediated debt in full. Quantum readiness belongs on the standard M&A due diligence checklist alongside conventional cybersecurity review: at minimum, whether a cryptographic inventory exists and its quality, the target's regulatory deadline exposure, and the estimated cost of remediation. Phase 1 already treats post-acquisition estates as a discovery problem; pre-close diligence is the opportunity to price that problem into the transaction instead of absorbing it afterward.

Audit and Assurance

Internal audit has a role in quantum security that most audit functions have not yet defined. The migration program is a multi-year transformation with significant budget, scope, and risk. It should be subject to the same audit oversight as any comparable enterprise program.

What to Audit

The cryptographic inventory (is it complete and current?). The migration program plan (is it realistic, resourced, and governed?). Migration execution (are systems actually migrated, or only marked as migrated?). The risk appetite statement (does it exist, is it board-approved, and are decisions consistent with it?). The KRI framework (are KRIs being measured, reported, and acted on?). Regulatory compliance (is the organization tracking the right requirements?). Vendor quantum readiness assessments (are they being conducted, and are findings being acted on?).

How to Evidence

The migration program should produce auditable artifacts at each stage: cryptographic inventory database with timestamps and coverage metrics; CBOM snapshots (quarterly); Steering Committee minutes showing KRI review and decision-making; risk appetite statement with board approval record; Regulatory Horizon Report with quarterly updates; vendor assessment records with remediation tracking; exception and deferral register with documented justifications and expiration dates; pilot test reports; training completion records.

Audit Timing

An initial audit within the first 6 months of the migration program (governance and inventory validation). A mid-program audit at 18–24 months (execution against plan). Annual audits thereafter until migration is substantially complete. The audit function

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

should also observe tabletop exercises (as described in the SOC Implementation section) to assess organizational incident preparedness.

Insurance Implications

Cyber insurance underwriters are incorporating quantum readiness into their risk assessments. Underwriters are asking about cryptographic inventories, PQC migration plans, and crypto-agility capabilities during policy renewals. Organizations unable to demonstrate a coherent transition plan may face higher premiums or explicit policy exclusions for quantum-related exposure.

The GRC function should ensure that the documentation produced by the migration program (the cryptographic inventory, the migration plan, the risk appetite statement, the KRI dashboard, the regulatory compliance tracker) is presentable to underwriters. An organization with a documented, governed, and measured migration program is a demonstrably better insurance risk than one with a vague intention to "migrate when standards are ready."

Crisis Communications and External Stakeholders

A quantum cryptographic event is a confidence crisis before it is a technical incident: the operational damage may take days to materialize, but the trust damage to customers, counterparties, regulators, and markets begins with the first headline. The SOC playbooks define the technical response; GRC owns the external one, and it cannot be drafted during the event.

Maintain pre-drafted communication templates as part of the GRC artifact set, one per audience: customer notification, regulator disclosure (mapped to each jurisdiction’s incident-reporting deadlines from the regulatory intelligence function), counterparty and partner notices for shared connections, and a market or investor statement for listed entities. Each template states what is known, what the organization had already migrated (the evidence dossier supplies the numbers) and what changes for the recipient. Name the spokesperson and the approval chain in advance, and exercise the communications leg in the same tabletop exercises that test the technical playbooks. Expect an attestation surge after any public quantum event: counterparties and customers will demand evidence of migration status, and the evidence dossier should be maintained so that answering is an export, not a project.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Industry Information Sharing

No organization sees enough of the quantum threat landscape alone. Participate in the sector's ISAC or equivalent peer network for PQC migration intelligence: vendor readiness experiences, interoperability findings, regulator expectations as they form, and early warning on implementation vulnerabilities. Share anonymized findings in both directions; the migration is ecosystem-wide, and organizations that learn from peers' pilots pay for fewer of their own mistakes. Treat ISAC participation as a standing CTI input alongside the feeds specified in the SOC Implementation foundation.

The GRC-SOC Handoff

The SOC Implementation section describes five detection use cases, each depending on the SOC having access to a queryable, up-to-date cryptographic posture registry. That registry is a GRC asset. The GRC-SOC handoff determines whether the detection capabilities the SOC builds actually function.

Beyond the posture registry, the GRC function produces four outputs that the SOC consumes:

The risk appetite statement, which defines the escalation thresholds for SOC alerts. A SOC analyst cannot determine whether a hybrid downgrade on a specific system warrants a Medium or High severity alert without knowing whether the data traversing that connection has HNDL-relevant secrecy requirements.

The regulatory compliance tracker, which tells the SOC which systems are highest priority for monitoring. Systems approaching a regulatory compliance deadline should receive more intensive monitoring.

The vendor readiness assessments, which inform the SOC's evaluation of third-party connections. A vendor connection where the vendor has no PQC timeline warrants different monitoring treatment.

The exception and deferral register, which the SOC must ingest as a dynamic suppression list. Without it, the SOC's cryptographic drift detection rules generate constant alerts from legacy systems that are legitimately deferred from migration, drowning out genuine regressions.

The relationship is symbiotic, but only if both sides invest in making their outputs consumable by the other. Designing these data flows should be a Phase 0 governance decision.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

GRC's Role in Tabletop Exercises

The GRC team should participate in every quantum-specific tabletop exercise described in the SOC Implementation section, but with a different focus. Where the SOC tests detection speed and response coordination, the GRC team tests governance:

In the **CRQC Announcement exercise,** the GRC function owns the board communication: who briefs the directors, what the message contains, what data supports it, and what decisions the board needs to make. If the GRC team discovers during the exercise that they cannot produce a board-ready exposure assessment within 4 hours because the cryptographic inventory is not queryable, that is the most valuable finding.

In the **Ambiguous Cryptanalysis Paper exercise,** the GRC function tests the decision framework: at what confidence level does an unverified cryptanalysis claim trigger action versus monitoring? If the risk appetite statement says "zero tolerance for algorithm compromise" but the Steering Committee's actual behavior during the exercise is "wait for peer review," the exercise exposes a gap between stated appetite and operational practice.

GRC Implementation Roadmap

Phase 1 (0–3 months). Draft the quantum risk appetite statement and present it to the board for approval. Register quantum risk formally in the enterprise risk register. Establish the five board-level KRIs and produce the first quarterly report. Begin integrating quantum readiness into the vendor assessment process for Tier 1 vendors.

Phase 2 (3–9 months). Implement the full executive and operational KRI framework. Establish the quarterly Regulatory Horizon Report. Conduct the initial internal audit of the migration program's governance and inventory. Begin preparing documentation for cyber insurance renewal.

Phase 3 (9–18 months). Operationalize the GRC-SOC handoff by making the cryptographic posture registry available to the SOC in machine-readable form. Integrate operational KRIs into the SOC dashboard. Incorporate quantum risk scenarios into the organization's existing enterprise risk scenario analysis. Expand vendor assessments to Tier 2 vendors.

Phase 4 (Ongoing). Refine KRI thresholds based on operational experience. Update the risk appetite statement annually. Maintain the regulatory intelligence process. Ensure audit findings are tracked and remediated.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

SECTOR ADAPTATION NOTES

The core methodology above applies universally. The following notes identify sector-specific constraints and priorities:

FINANCIAL SERVICES (PAYMENTS, BANKING)

- **Highest HNDL urgency:** Cross-border payment flows (SWIFT, correspondent banking) carry data with very long confidentiality requirements
- **Regulatory pressure:** PCI DSS v4.0 Req 12.3.3 (effective March 31, 2025) requires documentation and annual review of cryptographic cipher suites and protocols in use, effectively mandating cryptographic inventory for PCI-scoped entities; G7 Cyber Expert Group roadmap (January 2026); Europol QSFF prioritization framework; BIS Papers No. 158
- **Unique complexity:** Interbank payment systems involve ~320 cryptographic function calls for a single mobile banking session; cross-border payments involve 9+ parties and 30,000+ unique cryptographic functions
- **HSM dependency:** Payment HSMs (Thales payShield, Utimaco Atalla) have specific PQC upgrade timelines that constrain migration
- **Priority:** Key exchange on external-facing payment APIs; HSM-protected key material; SWIFT interface cryptography

TELECOMMUNICATIONS

- **Scale:** 5G networks involve cryptography at every layer from radio (5G-AKA) to core (service mesh, IMS, lawful intercept)
- **GSMA guidance:** The GSMA Post-Quantum Telco Network Task Force published the PQ.01–PQ.07 series: PQ.01 (impact assessment), PQ.02 (risk management), PQ.03 v2.0 (migration guidelines with Gantt charts for VPN, TLS, PKI, and MACsec migration, including CBOM guidance and crypto-agility methodology), and PQ.07 (non-terrestrial networks). This is the most operationally detailed sector-specific PQC guidance available.
- **Unique constraints:** Roaming interfaces, lawful intercept, non-terrestrial networks (satellite, HAPS), massive IoT device populations
- **Vendor concentration:** Small number of infrastructure vendors (Ericsson, Nokia, Huawei) control cryptographic implementations across the network core

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

- **Priority:** Backhaul encryption (site-to-site); 5G core service mesh; SIM/eSIM key management

CRITICAL INFRASTRUCTURE / OT

- **Longest asset lifecycles:** OT systems (SCADA, PLCs, RTUs) often have 15–25 year lifecycles with no update mechanism
- **Safety constraints:** Any change to OT cryptography must go through safety case re-certification in many jurisdictions
- **Minimal guidance:** CISA's October 2024 OT PQC guidance and RAND's 2023 study are in practice the only available frameworks
- **TNFL priority:** Firmware signing and safety certificate integrity are the primary quantum risks in OT: a forged firmware signature could compromise physical safety
- **Priority:** Gateway-based PQC termination at OT/IT boundary; firmware signing migration; long-term plan for device refresh

GOVERNMENT & DEFENSE

- **CNSA 2.0 compliance:** Mandatory and aggressively timed: 2027 for new NSS acquisitions, 2030 for software/firmware signing and networking equipment (VPNs/routers), 2033 for web/cloud/OS platforms, 2035 for all remaining NSS including custom and legacy systems. Note: SLH-DSA is not part of CNSA 2.0; ML-DSA-87 is the designated general CNSA 2.0 signature algorithm, with LMS/XMSS permitted for specific applications.
- **Classification constraints:** Working with classified systems adds complexity for testing and deployment
- **FedRAMP/FIPS requirements:** PQC implementations in US federal and defense systems must use FIPS 140-3 validated cryptographic modules. AWS-LC's FIPS module was the first publicly available FIPS 140-3 validated implementation to include ML-KEM support (albeit as a non-approved service pending CMVP algorithm list updates). Other validated modules from Thales, Utimaco, and others are expected to follow.
- **Priority:** Comply with CNSA 2.0 milestones; ensure PQC-capable products in procurement pipeline

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

APPENDICES

The appendices provide quick-reference tools, decision aids, and templates that support the framework's phases without interrupting the methodology narrative. They are designed to be printed, bookmarked, or extracted for use in workshops, steering committee meetings, and vendor engagements. Where an appendix relates to a specific phase, the relevant phase text includes a cross-reference, but the appendices are also usable as standalone reference cards for practitioners who are already familiar with the framework.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

APPENDIX A: ALGORITHM QUICK REFERENCE

NIST Standard	Algorithm	Type	Primary Use	Key/Signature Size	Notes
FIPS 203	**ML-KEM-768** (formerly CRYSTALS-Kyber)	Lattice-based KEM	Key exchange (TLS, VPN, SSH)	Public key: 1,184 bytes; Ciphertext: 1,088 bytes	Default recommendation for hybrid TLS (with X25519). ML-KEM-512 for constrained environments (IoT, embedded devices with limited RAM; ML-KEM-768 requires ~18 KB heap on typical MCUs); ML-KEM-1024 for highest security. On modern server hardware, performance overhead is negligible; on constrained devices, test memory and latency impact before selecting parameter set.
FIPS 204	**ML-DSA-65** (formerly CRYSTALS-Dilithium)	Lattice-based signature	Digital signatures (PKI, code signing, document signing)	Public key: 1,952 bytes; Signature: 3,309 bytes	Default recommendation for most signature use cases. ML-DSA-44 for constrained; ML-DSA-87 for highest security.
FIPS 205	**SLH-DSA** (formerly SPHINCS+)	Hash-based signature	Digital signatures (conservative choice)	Varies widely by parameter set	Stateless hash-based; no lattice assumptions. Larger signatures but different security foundation. Recommended as backup/alternative if lattice cryptanalysis advances.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Pending standardization (NIST):

- **FN-DSA (formerly FALCON):** Draft FIPS 206. Compact lattice-based signatures suitable for bandwidth-constrained environments, though it requires complex floating-point arithmetic to implement safely. Expected to complement ML-DSA for use cases where signature size is critical.
- **HQC:** Selected by NIST in March 2025 as an additional KEM to standardize, providing a non-lattice-based alternative to ML-KEM. Worth tracking as a diversification option even though it is not yet a final NIST standard.

Already standardized (NIST Special Publication):

- **HBS (LMS/XMSS):** NIST SP 800-208 (finalized October 2020). Stateful hash-based signature schemes approved for specific use cases such as firmware and software update signing, where cryptographic state can be managed safely. These are intended for niche, tightly controlled use cases rather than as general-purpose replacements for ML-DSA or SLH-DSA.

Regional variations:

- **BSI (Germany):** Recommends hybrid schemes with ML-KEM as primary; acknowledges FrodoKEM and Classic McEliece as having undergone extensive cryptanalysis (but does not elevate them to equal co-recommendations with ML-KEM)
- **South Korea (KpqC):** Sovereign algorithms SMAUG-T, NTRU+, HAETAE, AIMer (January 2025 final selections)

Standards Pipeline (June 2026 Update)

In addition to the three standardized algorithms above, organizations should track the following pipeline items for crypto-agility planning:

FN-DSA (FIPS 206) -- In Standardization. Based on the FALCON algorithm. Compact lattice-based digital signatures with smaller output than ML-DSA: approximately 666 bytes at NIST Security Level 1 versus 2,420 bytes for ML-DSA-44. Useful where signature size is a critical constraint. Draft in NIST/Commerce clearance as of June 2026; final standard expected late 2026 or early 2027. Do not deploy pre-standard. Plan for it in crypto-agility architecture. Systems that can accommodate FN-DSA when standardized will have more flexibility in signature-constrained environments.

HQC -- In Standardization. Code-based key encapsulation mechanism selected by NIST in March 2025 as an algorithmic diversity backup for ML-KEM. Uses a different

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

mathematical hardness assumption (error-correcting codes versus structured lattices). If a future attack compromises lattice-based KEMs, HQC provides a non-lattice fallback. Draft standardization in progress; final expected 2027. This is the strongest concrete argument for crypto-agility: organizations that hardcode ML-KEM without abstraction will face an emergency migration if lattice-based algorithms are weakened; organizations that build crypto-agile systems can add HQC as an alternative via configuration.

Additional Digital Signatures -- Third Round (9 candidates, May 2026). NIST announced on May 13, 2026 that nine candidates have advanced: FAEST, HAWK, MAYO, MQOM, QR-UOV, SDitH, SNOVA, SQIsign, and UOV. Evaluation expected to take approximately two years. These are not for production planning. They are relevant for one purpose: demonstrating to stakeholders why crypto-agility is not optional. The PQC algorithm set is still evolving, and any architecture that assumes ML-KEM and ML-DSA are permanent is making the same mistake that led to SHA-1-hardcoded systems and DES-locked hardware.

CNSA 2.0 Algorithm Requirements

For organizations subject to CNSA 2.0, algorithm selection is constrained. Key establishment requires ML-KEM-1024 (NIST Security Level 5) -- lower security levels are not permitted for national security systems. Digital signatures require ML-DSA-87 (NIST Security Level 5). SLH-DSA is excluded from CNSA 2.0 despite being a NIST-standardized algorithm -- NSA chose not to include it, likely due to its larger signature sizes and slower verification compared to ML-DSA. Symmetric cryptography remains AES-256 and hash functions remain SHA-384 or SHA-512, unchanged from CNSA 1.0.

Note that CNSA 2.0 is being adopted by financial services organizations as a "highest common denominator" standard even when not regulatorily required, because it simplifies compliance across government and commercial contexts.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

APPENDIX B: DECISION TREE — "WHERE DO I START?"

If you have done nothing yet: → Go to Phase 0. Secure executive mandate and budget. Nothing else works without this.

If you have budget but no inventory: → Go to Phase 1. Deploy automated discovery on your top 20 systems. Simultaneously begin Phase 2 CBOM structure design.

If you have an inventory but no migration plan: → Go to Phase 3. Score and prioritize your inventory. Then Phase 4 for roadmap.

If you have a plan but no pilots: → Go to Phase 5. Select your first hybrid TLS pilot on an internet-facing system you fully control.

If you have pilots running but infrastructure isn't ready to scale: → Go to Phase 6. Assess PKI, HSM, and network device readiness.

If you're blocked on vendors: → Go to Phase 7. Classify vendors by criticality, engage formally, deploy bridging patterns.

If you don't know where you stand: → Run the Maturity Model self-assessment (Cross-Cutting section). Your weakest domain tells you where to focus.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

APPENDIX C: MOSCA'S INEQUALITY — THE DECISION FRAMEWORK

Mosca's Theorem provides the mathematical framework for migration urgency:

If X + Y > Z, you must act now.

Where:

- **X** = Security shelf-life: the number of years the data you are protecting must remain confidential
- **Y** = Migration time: the number of years required to migrate your cryptographic infrastructure to quantum-safe
- **Z** = Threat timeline: the number of years until a cryptanalytically relevant quantum computer (CRQC) exists

Practical application:

- If your data needs 15 years of confidentiality (X=15) and migration takes 5 years (Y=5), then X+Y=20. If CRQC arrives in 15 years (Z=15), you are already 5 years too late.
- The PQCC roadmap distinguishes between "urgent adopters" (X+Y>Z) and "regular adopters" (X+Y≤Z) using this inequality.
- For HNDL-exposed data, X is the data's remaining confidentiality requirement. For TNFL-exposed signatures, X is the remaining validity period of the trust anchor.

The uncomfortable reality: For most large enterprises, Y (migration time) is 4–15 years. For sensitive data with long-lived confidentiality requirements, X (shelf-life) is 10–20 years. Most credible CRQC estimates place Z at 10–20 years. The math is uncomfortably tight, and it gets worse with every year of delay.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

APPENDIX D: HYBRID APPROACH JURISDICTIONAL COMPLIANCE MATRIX

Jurisdiction	Hybrid Position	Implication for Practitioners
BSI (Germany)	**Strongly recommended** for all PQC KEMs and signatures (except standalone hash-based signatures). Mandated in BSI certification ("visa") contexts.	Hybrid should be the default design assumption in German-facing deployments
ANSSI (France)	**Strongly emphasized** in the short and medium term; Phase 2 of ANSSI's timeline keeps hybrid mandatory for long-term-security claims in ANSSI's visa framework	Hybrid is the safer default for French assurance-sensitive products
Netherlands (AIVD/CWI/TNO)	**Recommended** in the PQC Migration Handbook, especially hybrid-AND internally; this is guidance, not a legal mandate	Strong guidance toward hybrid, with explicit policy for exceptions
NCSC UK	Prefers a **single migration to fully post-quantum PKI** rather than an intermediate hybrid PKI; acknowledges hybrid as an acceptable interim measure for confidentiality and interoperability	Use hybrid selectively, especially for key exchange / confidentiality; do not treat hybrid as the UK end-state
NIST (US)	Hybrid key establishment is **allowed but not required**; SP 800-227 discusses X-Wing as an example of a general-purpose hybrid KEM	Flexibility to choose hybrid or pure PQC depending on the system
NSA CNSA 2.0	**Accepted as interim** only	Hybrid acceptable during transition; end-state must be CNSA 2.0 compliant (pure PQC)
ASD (Australia)	**Not recommended, but not prohibited**	Pure PQC preferred in the long term; hybrid used only where interoperability or resiliency justifies it

Multinational recommendation: Deploy hybrid as the default deployment pattern. This satisfies the strictest requirements (continental European agency guidance), preserves interoperability during transition, and provides defense-in-depth against unknown PQC algorithm weaknesses. Plan the architecture for eventual transition to pure PQC to satisfy NCSC UK preferences and CNSA 2.0 end-state requirements. This single approach satisfies all jurisdictions simultaneously.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

APPENDIX E: QUICK-REFERENCE CHECKLISTS

90-Day Quick Start Checklist

1. ☐ Executive sponsor identified and committed
2. ☐ QRPM appointed
3. ☐ SteerCo convened; charter published; RACI defined
4. ☐ Multi-year budget (minimum 3 years) approved or in approval process
5. ☐ Top 20 critical systems identified (Phase 0 scoping)
6. ☐ Cryptographic discovery tool selected and procurement initiated
7. ☐ Quantum risk added to enterprise risk register
8. ☐ Cryptographic policy updated to include PQC-approved cipher suites
9. ☐ Procurement policy updated to require PQC roadmaps from vendors
10. ☐ Training cohort (10–20 people) enrolled in PQC fundamentals
11. ☐ First 2 hybrid pilot targets selected
12. ☐ Vendor questionnaires sent to top 10 strategic vendors
13. ☐ KPI baseline values established; Q+1 targets set
14. ☐ Board briefing delivered and documented

Quarterly Board Report Template

- **Overall maturity level:** [0–5] with trend arrow
- **Coverage KPI:** X% of Tier-1 endpoints on hybrid/PQC (target: Y%)
- **Inventory KPI:** X% of estate mapped in CBOM (target: Y%)
- **Vendor KPI:** X% of strategic vendors with dated PQC commitments (target: Y%)
- **Trust KPI:** X PKI/signing anchors shortened/rotated this quarter
- **Key risks:** [Top 3 risks with mitigation status]
- **Key decisions required:** [Any SteerCo escalations needing board input]
- **Budget status:** On track / Variance [±X%]
- **Next quarter milestones:** [3–5 concrete deliverables]

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

APPENDIX F: COMMON OBJECTIONS AND EVIDENCE-BASED RESPONSES

Program managers leading PQC migration will encounter recurring objections from stakeholders who resist commitment, budget, or urgency. This appendix provides evidence-based responses to the most common objections, drawn from practitioner experience across real migration programs. Each entry acknowledges the grain of truth in the objection before explaining why it is incomplete or incorrect. For a comprehensive treatment, see PostQuantum.com's "PQC Governance Objections: A Field Guide" and the forthcoming QuantumReady guide.

"Quantum computers are decades away." There is genuine uncertainty about CRQC timelines. But this objection confuses the quantum computing timeline with the migration timeline. Mosca's Inequality (X + Y > Z, where X is data sensitivity lifetime, Y is migration duration, and Z is time until CRQC) shows that organizations with long-lived sensitive data and multi-year migration timelines must start now regardless of the CRQC date. More importantly, HNDL is an active present-tense threat: adversaries are harvesting encrypted data today. Every day without hybrid key exchange creates permanent exposure. And regulatory deadlines (CNSA 2.0, NIS2 amendment, G7 roadmap) do not wait for Q-Day.

"Our vendors will handle it." Vendors will implement PQC in their products -- eventually. But vendor timelines are set by vendor business priorities, not yours. The framework's Phase 7 documents in detail why vendor delegation is dangerous: vendors control when PQC is available in their GA releases, not when your organization migrates. Many critical systems depend on multiple vendors, all of whom must ship PQC support before migration can proceed. Without proactive vendor engagement starting in Year 1, organizations discover vendor blocking issues in Year 3 when it is too late to influence vendor roadmaps.

"We will wait for standards to settle." The core standards are settled. NIST finalized FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), and FIPS 205 (SLH-DSA) in August 2024. Additional algorithms (FN-DSA, HQC) are in standardization but do not change the production-ready status of the primary standards. Organizations waiting for additional standards are deferring action that could begin today. The crypto-agility architecture recommended by this framework ensures that when additional algorithms are standardized, they can be adopted via configuration -- no architectural changes required.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

"We do not have budget." PQC migration does require investment. But framing it as a standalone budget line is strategically suboptimal. The infrastructure modernization umbrella strategy (see Phase 0, Activity 0.2c) positions PQC within broader IT modernization that is independently justified: PKI automation for shorter certificate lifetimes, FIPS 140-3 migration for the FIPS 140-2 sunset, HSM refresh for end-of-life hardware, and library upgrades for known CVE remediation. The PQC-specific incremental cost is a fraction of the total.

"This is just an IT problem." PQC migration touches every vendor relationship, every data flow, every compliance obligation, and every business process that depends on digital trust. The framework's Phase 0 governance structure intentionally includes business unit representatives, procurement, legal, and compliance -- because the migration scope extends far beyond the IT organization. Organizations that treat PQC as an IT project discover during Phase 5 that they lack the authority, budget, and cross-functional coordination to execute at enterprise scale.

"We will do it during the next refresh cycle." Technology refresh cycles are typically 3-7 years. Some embedded and OT systems have refresh cycles of 15-40 years. If the next refresh is 2030, the organization is planning to start PQC migration after multiple regulatory deadlines have already passed (CNSA 2.0 2030 deadline, EU critical infrastructure 2030 target, G7 financial critical systems 2030-2032). The framework's wave-based deployment model is specifically designed to avoid big-bang refresh dependencies: migrate what you can via software/configuration updates now, align hardware-dependent migrations with refresh cycles where possible, and deploy compensating controls (quantum-safe gateways, overlay encryption) for systems that cannot be migrated until refresh.

"Nobody in our sector has started." This is factually incorrect. F5 Labs measured 42% of the top 100 websites supporting hybrid PQC key exchange as early as mid-2025. Google, Cloudflare, Meta, Apple, and Signal have deployed PQC in production. Browser vendors have enabled PQC by default. OpenSSL, Go, and Node.js ship with PQC defaults. The G7 has published a PQC roadmap for financial services. The EU has proposed explicit PQC requirements in NIS2. Organizations that believe they are ahead of the curve by waiting are, in reality, falling behind an accelerating industry migration.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

"We did an inventory -- we are done." Inventory is Phase 1 of an 8-phase framework. Completing an inventory means the organization now knows what needs to be migrated. It does not mean anything has been migrated. The value of the inventory is entirely in what it enables: CBOM documentation (Phase 2), risk scoring and prioritization (Phase 3), roadmap construction (Phase 4), and migration execution (Phases 5-7). An inventory without a migration plan is an audit artifact, not a security outcome.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

APPENDIX G: CROSSWALK TO OTHER FRAMEWORKS

Organizations rarely adopt this framework in a vacuum: many have already committed to NIST CSF 2.0, the PQCC Migration Roadmap, ETSI TR 103 619, or the Dutch PQC Migration Handbook. The crosswalk below positions this framework as the execution layer under those commitments rather than a competitor to them. Mappings are indicative; the source documents differ in granularity.

Framework phase	**NIST CSF 2.0**	**PQCC Roadmap**	**ETSI TR 103 619**	**Dutch Handbook**
Phase 0 — Executive Mandate	Govern (GV.OC, GV.RM, GV.RR)	Preparation	Stage 1 (initiation)	Diagnosis
Phase 1 — Discovery & Inventory	Identify (ID.AM)	Baseline Understanding	Stage 1 (inventory compilation)	Diagnosis
Phase 2 — CBOM	Identify (ID.AM)	Baseline Understanding	Stage 1 (inventory compilation)	Diagnosis
Phase 3 — Risk Scoring	Identify (ID.RA)	Baseline Understanding	Stage 2 (migration planning)	Diagnosis / Planning
Phase 4 — Roadmap & Governance	Govern, Identify (GV.PO, ID.IM)	Planning and Execution	Stage 2 (migration planning)	Planning
Phase 5 — Pilots & Migration	Protect (PR.DS, PR.PS)	Planning and Execution	Stage 3 (migration execution)	Execution
Phase 6 — Infrastructure	Protect (PR.PS, PR.IR)	Planning and Execution	Stage 3 (migration execution)	Execution

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Phase 7 — Vendor & Supply Chain	Govern (GV.SC)	Planning and Execution	Stages 2–3 (supplier actions)	Planning / Execution
Program Foundations (SOC, GRC, agility, skills, metrics)	Detect, Respond, Recover; Govern	Monitoring and Evaluation	Cross-stage	Cross-step

Milestone-based guidance maps onto the lifecycle rather than onto individual phases. NCSC UK's milestones align as follows: discovery and planning complete by 2028 corresponds to Phases 0 through 4; high-priority migration by 2031 to the first waves of Phases 5 and 6; full migration by 2035 to wave completion plus the Migration Verification & Program Closure protocol. CNSA 2.0's milestones gate specific activities: the 2027 acquisition requirement lands in Phase 7 procurement language, the 2030 software and firmware signing deadline in Track B's SP 800-208 deployment, the 2030 networking equipment deadline in Track A's infrastructure waves, and the 2033 and 2035 horizons in the outer years of the multi-year roadmap.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

APPENDIX H: PROTOCOL COVERAGE MATRIX

Use this matrix as a completeness check against your own estate: every protocol family below carries quantum-vulnerable cryptography, and several are routinely forgotten because no single team owns them. Deployability reflects the state of standards and implementations as of June 2026 and changes quickly; verify before planning. Framework references point to where each family is addressed.

Protocol family	**Quantum threat**	**Standards status (June 2026)**	**Deployability**	**Framework reference**
TLS 1.3 (web, APIs)	HNDL (key exchange); TNFL (certificates)	Hybrid ML-KEM deployed at scale; MTC drafts for public Web PKI	Deploy now (Track A); plan the PKI fork	Activity 5.2; Phase 6 PKI
VPN / IPsec (IKEv2)	HNDL	RFC 9370 hybrid IKEv2	Deploy now	Activity 5.2
SSH	HNDL; TNFL (host and user keys)	Hybrid default in current OpenSSH	Deploy now	Activity 5.2
Email transport (SMTP over TLS)	HNDL	Same hybrid TLS mechanisms	Deploy as mail servers support it	Activity 5.2
Email signing (S/MIME, PGP)	TNFL	IETF LAMPS drafts; immature	Monitor; shorten key lifetimes	Track B
DKIM	TNFL (RSA throughout)	No PQC standard yet	Monitor; rotate keys; track IETF	Track B
DNSSEC	TNFL only	IETF strategy work; ML-DSA/SLH-DSA research; root KSK rollover measured in years	Monitor only; no unilateral action possible	Track B

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Code and firmware signing	TNFL (highest persistence)	SP 800-208 LMS/XMSS final; composite ML-DSA drafts	Deploy now (LMS/XMSS, dual-signed)	Track B; Activity 5.2
Document signing and RFC 3161 time-stamping	TNFL (decades-long validity)	PQC time-stamping immature; ERS (RFC 4998) re-anchoring patterns	Plan re-signing and re-timestamping cycles for archives	Track B; Closure
Long-lived signed records and archives	TNFL	Re-sign versus Merkle-anchor strategies; evidence record syntax	Decide per archive; document in the CBOM	Activity 5.6; Closure
Kerberos / Active Directory (PKINIT)	TNFL; HNDL on ticket exchange	Vendor-dependent; early	Track operating-system vendor roadmaps	Track B identity
FIDO2 / WebAuthn passkeys	TNFL (ES256/RS256 credentials)	PQC COSE codepoints assigned; ML-DSA drafts active	Monitor; plan re-enrollment; ask CIAM vendors	Track B identity
OIDC / SAML token signing	TNFL	Algorithm-agile by design; PQC profiles early	Inventory signing keys; migrate with IdP support	Track B identity
802.1X EAP-TLS (network access)	HNDL; TNFL (certificates)	Follows TLS and internal PKI	Migrate with internal PKI	Phase 6 PKI
Wi-Fi WPA3 (enterprise)	Enterprise mode rests on EAP-TLS; SAE is symmetric	Follows EAP-TLS	Cover via the 802.1X row	Phase 6

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Constrained IoT / LPWAN / satellite	HNDL; TNFL (device identity)	Profiles immature; severely size-constrained	Gateway termination; PSK overlays; see sector extensions	Activities 6.3, 5.5

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ABOUT THIS VERSION

This framework is a living document, versioned and updated as standards evolve, vendor products mature, regulatory deadlines shift, and practitioner experience accumulates across real migration programs. Version 2.1 was published in June 2026 as a targeted update to the June 2026 v2.0 release, which itself updated the March 2026 v1.1 to reflect significant developments in PQC deployment reality.

WHAT'S NEW IN V2.1

Version 2.1 is a targeted update, not an architectural revision. The 8-phase structure, the two-track model, and the v2.0 planning assumptions are unchanged. The update does five things: it makes the two-track model fully consistent throughout the document, takes two new positions the deployment environment now demands, ports proven material from the source article library into the methodology, adds the closing protocol the lifecycle previously lacked, and extends the methodology into execution domains it previously left implicit: AI-assisted tooling, data at rest, counterparties, cloud and SaaS, and alignment with the frameworks organizations have already adopted.

Consistency with the two-track model. Activity 3.3 no longer presents key exchange, signatures, and data at rest as one serial priority list. Sequencing now operates within Track A and Track B as parallel workstreams, with the cascade-asymmetry rationale for starting signature migration early despite TNFL's lower urgency. The illustrative risk appetite statements were also revised to dates achievable under the deployment environment classification.

Two new positions. The framework now takes an explicit position on hybrid and composite signatures (default to composite or dual-signing; solo ML-DSA is a recorded risk decision, not a default; see Phase 5), and introduces algorithm-specific vulnerability weighting in Phase 3: quantum attack cost tracks key size, not classical strength, so ECC-protected estates must not be deprioritized relative to RSA.

New protective and closing content. Activity 2.5 (Secure the CBOM and Program Artifacts) treats the program's own outputs as the high-value intelligence targets they are. A new Migration Verification & Program Closure section defines the evidence standard for declaring systems migrated, the decommissioning of classical material, and the formal transition to business-as-usual.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Scope and method additions. The identity and authentication stack (FIDO2/WebAuthn, token signing, Kerberos PKINIT, 802.1X) is now explicitly in Track B scope; SP 800-208 stateful hash-based signatures are foregrounded as the deploy-now component of Track B; Phase 0 gains reference program economics drawn from a published full-program description; Phase 1 gains a confidentiality-horizon determination method; the evidence dossier's litigation-defense purpose is made explicit; and M&A due diligence, procurement model-language references, and additional common failures round out the update. The web PQC adoption datapoint is corrected to its F5 Labs mid-2025 attribution, and cross-references to the companion book, Quantum Ready, have been added.

Execution and ecosystem additions. Phase 5 gains a data-at-rest decision framework (Activity 5.6) and an explicit position on AI-assisted migration (Activity 5.7, with a matching AI tool category in Phase 1). Phase 7 gains a counterparty coordination pattern for parties the organization cannot contractually compel (Activity 7.6) and a cloud shared-responsibility and SaaS treatment (Activity 7.7). Phase 4 gains a pre-drafted Accelerated Execution Profile (Activity 4.7). The Metrics section adds a risk-weighted coverage companion KPI; the regulatory map adds algorithm sovereignty and standards fragmentation; the GRC foundation adds crisis communications and industry information sharing; the Skills foundation adds a skills matrix; and How to Use This Framework adds role-based reading paths and right-sizing profiles. Two new appendices close the update: a crosswalk to NIST CSF 2.0, the PQCC Roadmap, ETSI TR 103 619, and the Dutch Handbook (Appendix G), and a protocol coverage matrix spanning TLS to DNSSEC, DKIM, time-stamping, and constrained IoT (Appendix H).

WHAT'S NEW IN V2.0

Version 2.0 retains the same 8-phase architecture established in v1.1. What has changed is the planning model beneath it -- reflecting a quarter in which PQC migration moved from standards readiness to deployment reality. Three methodological changes drove this major version:

Two-track migration model. Key exchange migration (to stop HNDL exposure) and signature/authentication migration (to address TNFL and prepare for PKI evolution) are now explicitly treated as parallel tracks with different urgency drivers, deployment patterns, and infrastructure dependencies. See Phase 5.

PKI architecture fork. Google’s Merkle Tree Certificate architecture, backed by Cloudflare and Let’s Encrypt, means public Web PKI is heading toward a different trust model.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

Organizations must now classify their PKI deployments and plan for two parallel architectures. See Phase 6.

FIPS validation gap. No FIPS 140-3 validated cryptographic module offers PQC algorithms in approved mode as of June 2026. For regulated environments, this gap determines when PQC can enter production. A new deployment environment classification guides roadmap construction. See Phase 5.

Beyond these three methodological changes, v2.0 substantially expands the Program Foundations with two new sections and three major expansions.

SOC Implementation specifies the Security Operations Center's role in PQC migration: five detection use cases (hybrid downgrade, cryptographic drift, certificate lifecycle anomalies, TNFL/signature integrity monitoring, enhanced HNDL-indicator detection) with illustrative rules and thresholds, a three-horizon quantum cyber threat intelligence model (tactical, operational, strategic), four incident response playbooks (algorithm vulnerability disclosure, confirmed hybrid downgrade attack, credible CRQC announcement, emergency algorithm rotation), five tabletop exercise scenarios, and a phased SOC implementation roadmap.

GRC Implementation specifies the Governance, Risk, and Compliance function's role: a 17-indicator Key Risk Indicator framework cascading across three organizational levels (board, CISO, operational) with illustrative thresholds, quantum risk appetite statement templates with operational tolerances, a regulatory intelligence process (quarterly Regulatory Horizon Report), third-party quantum readiness assessment methodology, audit and assurance procedures, insurance preparation guidance, and the GRC-SOC handoff that makes detection capabilities function.

Governance has been expanded from a reference paragraph to a full subsection covering program leadership (who should own the program and why), board oversight (minimum engagement model with quarterly KPIs and escalation triggers), the quantum risk appetite statement, and three lines of defense mapping for PQC migration.

Crypto-Agility has been expanded from architectural principles to a five-dimensional operational discipline (architecture, operations, governance, skills, supply chain), each with a testable criterion, a four-year implementation roadmap integrated into the migration phases, and six OKRs with measurement methods.

Skills & Team Structure has been expanded with team sizing heuristics, a build/borrow/buy decision matrix for six capability areas, a four-level training approach, and the crypto champion program.

Version 2.0 also incorporates: an updated regulatory timeline reflecting 14+ deadlines converging between 2026 and 2030; expanded cost estimation methodology; NIST CSWP

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

39 crypto-agility alignment; algorithm pipeline updates (FN-DSA, HQC, 9 additional signature candidates); current PQC deployment ecosystem data; a multinational regulatory navigation framework; and an objection-handling reference appendix (Appendix F). All sector extensions have been updated to v2.0.

CURRENCY OF TECHNICAL REFERENCES

Vendor capabilities, HSM firmware versions, MTC/PLANTS standardization status, and tool categories cited throughout this framework reflect the market as of the publication date. PQC vendor capabilities and standards status change frequently; readers should verify current GA dates, FIPS 140-3 validation status, and product capabilities against original vendor sources or PostQuantum.com, which tracks these changes continuously.

STANDARDS IN PROGRESS

NIST IR 8547 (Initial Public Draft November 2024; final publication expected 2026) sets the deprecation and disallowance timelines referenced throughout this framework. Federal agencies and their contractors should reference the final published version, and the application-specific guidance it supersedes (SP 800-52 for TLS, SP 800-77 for IPsec, SP 800-175B for general cryptographic standards), as these documents are updated. Similarly, CNSA 2.0 milestone dates, ETSI technical reports, and sector-specific regulatory guidance (PCI DSS, DORA, GSMA PQ series) may be updated between framework versions.

ENGAGEMENT

For questions, advisory engagement, or sector-specific framework adaptation, contact Marin Ivezic via PostQuantum.com or AppliedQuantum.com. Organizations seeking structured support for any framework phase, from executive business case development through migration execution, can engage Applied Quantum's advisory practice directly at appliedquantum.com.

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

ABOUT

ABOUT THE AUTHOR

Marin Ivezic is a former Fortune Global 500 CISO/CTO, the Editor-in-Chief of PostQuantum.com, and the founder and CEO of Applied Quantum. He previously held cybersecurity partner and practice lead roles, including regional and global leadership positions, at IBM, Accenture, PwC, and KPMG. He is also the former CEO of Cyber Agency and the founder of Cryptosec, a cryptography advisory and engineering firm.

Marin has been involved in quantum technologies for over 25 years, having advised governments on the quantum threat since the early 2000s. He previously founded Boston Photonics and PQ Defense. He has led real-world quantum readiness programs for telecoms, financial institutions, and critical infrastructure operators, including programs exceeding 120,000 discrete migration tasks.

PostQuantum.com, his personal blog, reaches over 1 million monthly readers and is recognized as the premier quantum security publication for CISOs, CTOs, and security architects worldwide.

QUANTUM READY: THE COMPANION BOOK

Quantum Ready (QuantumReady.com) is the book-length companion to this framework, written by the same author. Where this document is deliberately methodology-grade (prerequisites, activities, outputs, decision logic). The book provides the complete treatment: the reasoning behind each phase, extended case examples from real migration programs, sector narratives, and guidance for leading the program from the first board conversation through closure. The two are maintained in alignment: the framework is updated as the field moves, and the book supplies the depth that a methodology document omits by design.

ABOUT APPLIED QUANTUM

Applied Quantum (AppliedQuantum.com) is a research-driven professional services firm focused entirely on quantum technologies and post-quantum security. Its dedicated security practice, Secure Quantum (SecureQuantum.com), focuses on quantum security

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.

advisory and PQC migration. Services span Quantum Security & PQC Migration, Quantum Strategy & Sovereignty, Quantum Engineering & Implementation, and Quantum Commercialization.

If you need help: Applied Quantum provides advisory, program design, and hands-on migration execution support. Contact: contact@appliedquantum.com

© 2026 Marin Ivezic / Applied Quantum. Licensed under CC BY 4.0.