SEPTEMBER 2026

APPLIED QUANTUM

THE APPLIED QUANTUM PQC MIGRATION FRAMEWORK

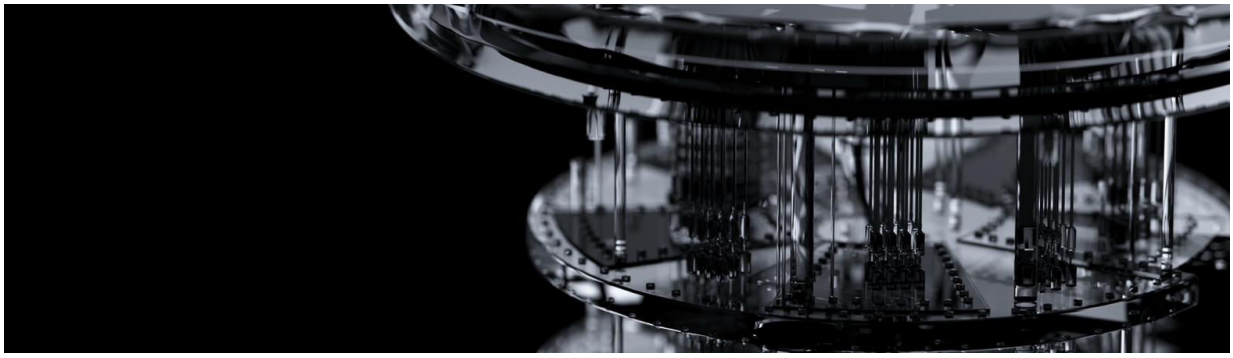

From Discovery to Deployment – *the phase-by-phase guide to migrating enterprise cryptography before quantum computers arrive*

Version 3.0 – September 2026

Marin Ivezic

CEO, Applied Quantum

Author, PostQuantum.com

Disclosure: Applied Quantum provides post-quantum migration advisory services. This work is published openly under CC BY 4.0 and is usable without engaging Applied Quantum.

"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 September 2026. These references may become outdated quickly. Readers should verify current status against primary sources before making implementation decisions; movement between editions is published through the PQC Migration Brief.

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, remains an Initial Public Draft at the date of this edition; the deprecation and disallowance dates it proposes are cited as anchors and re-checked at each release. Federal agencies and their contractors should reference the final version when it publishes.

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) and the crypto-agility operating model. Sector-specific adaptations included as summary notes; dedicated sector extensions published separately.
Status	Version 3.0 - published September 2026

VERSION HISTORY

Version	Date	Changes
1.0	June 2025	Initial publication. 8-phase lifecycle, cross-cutting sections, 5 appendices
1.1	March 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	June 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 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

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

Version	Date	Changes
		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).
3.0	September 2026	Major version. Crypto-agility becomes the program's organizing method and measured end state (dual-outcome charter, a funded Crypto-Agility workstream, the agility criterion at every gate, per-wave rehearsals, the closure proof, the agility-debt register). Plain-standard register with a normative convention; Applying This Phase, closing common-failure lines and Challenge Questions in every phase; worked artifacts in Appendix I. Own minimum CBOM record with the Applied Quantum CBOM Profile cited by minimum version and the Cryptographic Concentration Framework cross-referenced; evidence grades and a coverage denominator with a signed accepted-gaps register. Corrections from v2.x: HNDL restated as inference from demonstrated capability, with the CISA/NSA/NIST, AIVD/CWI/TNO and Federal Reserve staff statements quoted in Appendix F; SOC Use Case 5 reframed as exfiltration detection; the v2.0 FIPS validation-gap statement corrected (CMVP #5247); CNSSP 15 milestones added. Standards currency through September 2026 (RFC 9954, 10024, 9964, 10042, 9980, 9881, 9882; the 2026 FIPS 140-3 IG updates; CycloneDX 1.7 / ECMA-424; FN-DSA and HQC status; EO 14412's CBOM-guidance and FAR directions). Payments and Digital Assets sector notes. Method decisions adopted after review: the scoring arithmetic with unknown handling; a delivery state on every scored entry; two results per gate record; closure at zero unaccepted agility debt; rollout, wave-exit and closure gates; the three lines of defense by function;

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

Version	Date	Changes
		one maturity scale across phases; a requirements register generated from the text (Appendix J); Marin's Law as a design heuristic. The full account is in What's New in v3.0.

HOW TO CITE THIS FRAMEWORK

Ivezic, M. (2026). The Applied Quantum PQC Migration Framework – Universal, Version 3.0. Applied Quantum. https://pqcframework.org

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. Reusers may hold copyright in their own new contributions, but may not apply terms or technological measures that restrict others from exercising the rights the license grants in this material. Full attribution requirements and the framework's publication history are at PQCFramework.org/license.

RELATIONSHIP TO OTHER GUIDANCE

This framework was first drafted in March 2023 (v0.1) and first made publicly available under CC BY 4.0 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.org/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.07 (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

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

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

The body of this edition names categories and capability requirements. Commercial vendors appear in the body only where a dated public fact is itself the evidence (an announcement, a measurement, a published roadmap); open-source libraries and kits are named where a practitioner needs them. The references below carry the product names the body relied on in earlier editions, as named in the June 2026 edition and verified September 2026, and are maintained on PQCFramework.org between editions.

Dedicated cryptographic discovery platforms (Phase 1, Category 1): SandboxAQ, IBM Quantum Safe, Keyfactor (including InfoSec Global and CipherInsights), ISARA Advance, CryptoNext, Palo Alto Networks, and emerging specialists; see the vendor analysis on PostQuantum.com for capability comparisons.

General-purpose HSMs (Activity 6.2): Thales Luna HSM firmware 7.9 (July 2025) added native ML-KEM and ML-DSA, with FIPS 140-3 Level 3 validation of that firmware in progress at release; earlier 7.8 firmware did not carry the lattice algorithms. Utimaco Quantum Protect (April 2025) is an application package for the u.trust General Purpose HSM Se-Series, activated in the field without a hardware exchange, carrying ML-KEM, ML-DSA, LMS/HSS and XMSS/XMSS-MT; older Utimaco models outside the Se-Series are a hardware question. Payment HSMs (Sector Adaptation Notes): the Thales payShield and Utimaco Atalla product lines carry their own PQC firmware and PCI certification timelines, unverified at this edition.

Cloud key management (Activity 6.2): AWS KMS, Azure Key Vault and Google Cloud KMS publish their own PQC key-type roadmaps; check the current status per service and region.

Validated modules (Phase 5): CMVP certificate #5247 (Go Cryptographic Module, FIPS 140-3 Level 1 software, April 2026) carries ML-KEM key generation and encapsulation-decapsulation at ML-KEM-768 and ML-KEM-1024 among its approved algorithms; CMVP certificate #5497 (Crypto4A QASM Cryptographic Module, FIPS 140-3 Level 3 hardware, August 2026) carries ML-KEM, ML-DSA, SLH-DSA and LMS among its approved algorithms, with HSS vendor-affirmed rather than CAVP-tested and used alongside ECDSA in the firmware-update verification service of that module, and with the Security Policy qualifications stated in Phase 5. Other modules list PQC only as non-approved services. Consult the CMVP validated-modules list for the current state.

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

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

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.org, a dedicated companion site for this framework, drawing on relevant content from PostQuantum.com.

Two sibling Applied Quantum properties carry work this framework cites rather than repeats. The Applied Quantum CBOM Profile, the property taxonomy layered on CycloneDX that Phase 2 cites by minimum version, publishes at CBOMProfile.org. The Cryptographic Concentration Framework, which measures cryptographic concentration beneath the vendor layer and consumes a Profile-conforming CBOM, publishes at CCFramework.org. Both are open under CC BY 4.0 and follow this framework's sector taxonomy.

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). Quantum Academy (QuantumAcademy.com), founded by the author, offers trainings and certifications on the topics this framework covers, for the roles and training levels described in Skills & Team Structure.

The PQC Migration Brief (pqcmigrationbrief.com) carries corrections, standards and regulatory movement between editions, and the current jurisdiction roster is maintained on the Global PQC Migration Clock page on PostQuantum.com; the framework itself remains free and ungated.

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

TABLE OF CONTENTS

How to Use This Framework 18
Purpose and Use 18
Normative Convention 19
Three Layer Systems 20
Entry Points by Reader 20
Right-Sizing the Framework 21
Keeping Current 23
Framework Architecture at a Glance 23
Regulatory Timeline Context 24
Phase 0 – Executive Mandate & Business Case 29
Purpose 29
Parallelization note 30
Prerequisites 30
Framework prerequisites (from earlier phases) 30
Organizational prerequisites 30
Activities 31
0.1 Frame the Business Case 31
0.2 Build the Budget Structure 32
0.2b Build the Business Case – Additional Benefit Arguments 33
0.2c Develop Migration Cost Estimates 34
0.3 Establish Governance Structure 37
0.4 Draft the Program Charter 42
0.5 Conduct Initial Scoping Assessment 43
Applying This Phase 43
Outputs 44
Interdependencies 44
Backward dependencies 44
Feeds into 44
Runs in parallel with 45
Common Failures 45
Challenge Questions 46

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

Maturity Indicators 46
Phase 1 – Discovery & Inventory 47
Purpose 47
Parallelization note 47
Prerequisites 48
Framework prerequisites (from earlier phases) 48
Organizational prerequisites 48
Activities 49
1.0 Risk-Driven Scoping – Decide What to Inventory First 49
1.1 Establish Three Parallel Inventories 50
1.2 Deploy Cryptographic Discovery – Layered Approach 52
1.3 Map the Cryptographic Estate 57
1.4 Address the Asset Discovery Problem 58
1.5 Integrate with Existing Data Sources 61
1.6 Establish Continuous Discovery 62
Applying This Phase 66
Outputs 67
Interdependencies 68
Backward dependencies 68
Feeds into 68
Runs in parallel with 69
Common Failures 69
Challenge Questions 70
Maturity Indicators 70
Phase 2 – CBOM & Documentation 72
Purpose 72
Parallelization note 72
Prerequisites 73
Framework prerequisites (from earlier phases) 73
Organizational prerequisites 73
The Minimum Viable CBOM Model 73
Activities 74

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

2.1 Select CBOM Format and Tooling 74
2.2 Populate CBOM from Inventory Data 79
2.3 Integrate CBOM into Operational Processes 79
2.4 Establish CBOM Freshness Governance 80
2.5 Secure the CBOM and Program Artifacts 80
Applying This Phase 83
Outputs 83
Interdependencies 84
Backward dependencies 84
Feeds into 84
Runs in parallel with 85
Common Failures 85
Challenge Questions 85
Maturity Indicators 85
Phase 3 – Risk Scoring & Prioritization 87
Purpose 87
Parallelization note 87
Prerequisites 88
Framework prerequisites (from earlier phases) 88
Organizational prerequisites 88
Activities 89
3.1 Define Risk Scoring Model 89
3.2 Calculate Priority Scores 96
3.3 Apply Migration Sequencing Logic 99
3.4 Produce the Quantum Readiness Assessment (QRA) 100
Applying This Phase 100
Outputs 101
Interdependencies 102
Backward dependencies 102
Feeds into 102
Runs in parallel with 102
Common Failures 103

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

Challenge Questions 103
Maturity Indicators 104
Phase 4 – Roadmap & Governance 105
Purpose 105
Parallelization note 105
Prerequisites 106
Framework prerequisites (from earlier phases) 106
Organizational prerequisites 106
Activities 107
4.1 Define Year-1 Starter Plan (90-Day Governance Sprint) 107
4.2 Structure the Multi-Year Roadmap 109
4.3 Align to Infrastructure Refresh Cycles 110
4.4 Establish PMO Structure for Scale 111
4.5 Manage the Roadmap as a Living Instrument 113
4.6 Define Milestone Gates 114
4.7 Pre-Draft the Accelerated Execution Profile 117
4.8 Re-Baseline Protocol 117
Applying This Phase 118
Outputs 119
Interdependencies 119
Backward dependencies 119
Feeds into 119
Runs in parallel with 120
Common Failures 120
Challenge Questions 121
Maturity Indicators 121
Phase 5 – Pilots & Migration Execution 123
Purpose 123
Parallelization note 123
Prerequisites 124
Framework prerequisites (from earlier phases) 124
Organizational prerequisites 124

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

Deployment Environment Classification 125
Understanding Deployment Constraints 125
Two-Track Migration Model 130
Track A (Confidentiality) and Track B (Integrity and Authentication) 130
PQC Deployment Status (data as of June 2026) 134
Activities 135
5.1 Select Pilot Targets 135
5.2 Design Hybrid Deployments 135
5.3 Execute Pilots with Measurement 143
5.4 Scale from Pilot to Production Through Waves 143
5.5 Implement Defense-in-Depth Beyond Pure PQC 144
5.6 Decide the Data-at-Rest Strategy 146
5.7 AI-Assisted Migration Under the Existing Review Gate 148
Applying This Phase 148
Outputs 149
Interdependencies 150
Backward dependencies 150
Feeds into 150
Runs in parallel with 151
Common Failures 151
Challenge Questions 152
Maturity Indicators 152
Phase 6 – Infrastructure Modernization & Performance 154
Purpose 154
Parallelization note 154
Prerequisites 155
Framework prerequisites (from earlier phases) 155
Organizational prerequisites 155
PKI Architecture Evolution 156
The PKI Architecture Fork 156
Activities 158
6.1 PKI Modernization 158

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

6.2 HSM and KMS Modernization 160
6.3 Network Infrastructure Assessment 162
6.4 Performance Testing Methodology 164
6.5 Capacity Planning for PQC at Scale 164
Applying This Phase 165
Outputs 165
Interdependencies 166
Backward dependencies 166
Feeds into 166
Runs in parallel with 166
Common Failures 167
Challenge Questions 168
Maturity Indicators 168
Phase 7 – Vendor & Supply Chain Governance 170
Purpose 170
Parallelization note 170
Prerequisites 171
Framework prerequisites (from earlier phases) 171
Organizational prerequisites 171
Activities 172
7.1 Classify Vendor Portfolio by PQC Impact 172
7.2 Execute Vendor Engagement 172
7.3 Insert PQC Requirements into Procurement 174
7.4 Manage Vendor-as-Blocker Scenarios 176
7.5 Establish Ongoing Vendor Governance 176
7.6 Coordinate Counterparties You Cannot Contractually Compel 177
7.7 Cloud Shared Responsibility and the SaaS Class 178
Applying This Phase 179
Outputs 179
Interdependencies 180
Backward dependencies 180
Feeds into 180

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

Runs in parallel with 180
Common Failures 181
Challenge Questions 181
Maturity Indicators 182
Migration Verification & Program Closure 183
Verification Evidence Standard 183
Evidence by Cryptographic Function 184
Decommissioning Classical Material 187
Program Closure and Transition to BAU 187
Outputs 189
Common Failures 189
The Seven Program Foundations 191
Maturity Levels 191
Assessment Across Seven Domains 192
Self-Assessment Scoring 194
Metrics, KPIs & Reporting 194
Board-Level KPI Pack (Report Quarterly) 194
Operational KPIs (Report Monthly to SteerCo) 197
Evidence Dossier (for Audit and Regulatory) 197
Crypto-Agility as End-State Architecture 198
Why Crypto-Agility, Not Just PQC 198
The Five Dimensions of Crypto-Agility 199
Crypto-Agility Architecture Principles 200
Mechanisms by Layer 201
The Rehearsal Method 203
Crypto-Agility Implementation Roadmap 203
Crypto-Agility OKRs 204
Alignment with NIST CSWP 39 204
Regulatory & Standards Alignment Map 205
Mapping Framework Phases to Regulatory Requirements 205
Deadline Archetypes 207
Enforcement Tiers 208

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

Regulatory Pattern Since v2.1 ... 209
Regulatory Developments Recorded at v2.0 and v2.1 (history) ... 209
Multinational Regulatory Navigation ... 210
Algorithm Sovereignty and Standards Fragmentation ... 211
Skills & Team Structure ... 211
The Skills Challenge ... 211
Core Roles ... 212
Team Sizing ... 213
Skills Matrix ... 213
Build, Borrow, or Buy ... 214
Training Approach ... 216
Sustaining Capability Beyond the Migration ... 216
SOC Implementation ... 217
Prerequisites ... 217
Detection Use Cases ... 218
Cyber Threat Intelligence ... 222
Incident Response Playbooks ... 225
Tabletop Exercises ... 227
SOC Metrics Summary ... 228
Skills and Tooling Gaps ... 229
SOC Implementation Roadmap ... 229
GRC Implementation ... 229
Why Quantum Risk Breaks the Standard ERM Playbook ... 230
Risk Appetite and Tolerance ... 230
Key Risk Indicators at Board, CISO and Operational Levels ... 231
Regulatory Intelligence ... 234
Third-Party Quantum Readiness Assessment ... 235
Audit and Assurance ... 236
Insurance Implications ... 238
Qualified Trust Services and Electronic Signature Law ... 238
Crisis Communications and External Stakeholders ... 240
Industry Information Sharing ... 240

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

The GRC-SOC Handoff 240
GRC's Role in Tabletop Exercises 241
GRC Implementation Roadmap 241
Sector Adaptation Notes 243
Financial Services (Banking, Capital Markets, Insurance) 243
Telecommunications 243
Critical Infrastructure / OT 244
Government & Defense 244
Payments 245
Digital Assets 246
Appendices 247
Appendix A: Algorithm Quick Reference 248
Standards Pipeline (September 2026) 250
CNSA 2.0 Algorithm Requirements 250
Appendix B: Decision Tree – "Where Do I Start?" 252
Appendix C: Mosca's Inequality – The Decision Framework 253
Appendix D: Hybrid Approach Jurisdictional Compliance Matrix 254
Appendix E: Quick-Reference Checklists 257
90-Day Quick Start Checklist 257
Quarterly Board Report Template 257
Appendix F: Common Objections and Evidence-Based Responses 258
Appendix G: Crosswalk to Other Frameworks 262
Supervisory Readiness Indices 263
Sibling Instruments 264
Crypto-Agility Maturity Models 265
Appendix H: Protocol Coverage Matrix 266
Appendix I: Worked Examples Across the Phases 271
I.1 Phase 0 – Charter Purpose Statement in the Dual-Outcome Form 272
I.2 Phase 1 – Priority A Scoping Decision 272
I.3 Phases 2 and 3 – One Scored CBOM Entry 273
I.4 Phase 4 – Gate Record with the Agility Criterion 274
I.5 Phase 5 – Wave Rehearsal Record 275

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

I.6 Phase 6 – HSM Register Row 275
I.7 Phase 7 – Questionnaire Response Graded for Firmness 276
I.8 Phase 4 – Re-Baseline Protocol Trigger Record 277
I.9 Phase 3 – Appeals Record under the Ratified Rule 277
I.10 Closure – Decision Record 278
Appendix J: Requirements Register 280
About This Version 290
What's New in v3.0 290
What's New in v2.1 292
What's New in v2.0 293
Currency of technical references 295
Standards in progress 295
Engagement 295
About 296
About the Author 296
Quantum Ready: The Companion Book 296
Quantum Academy 296
About Applied Quantum 297

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

HOW TO USE THIS FRAMEWORK

PURPOSE AND USE

This document defines an executable methodology for enterprise post-quantum cryptography migration: eight phases (Phase 0, Executive Mandate, through Phase 7, Vendor & Supply Chain Governance), the program foundations that span them (maturity model, metrics, crypto-agility, regulatory alignment, skills, SOC and GRC implementation), a verification and closure protocol, and reference appendices. The program's mandate is readiness against the organization's dated obligations. The organization meets that mandate by building crypto-agility. The PQC migration is the first exercise of that capability. The program ends only when the organization can change cryptography on a planned date, and when a named owner keeps that capability running after the program closes.

Four areas are out of scope and covered by sibling documents:

- **Cryptographic concentration measurement** beneath the vendor layer is defined by the Cryptographic Concentration Framework (CCFramework.org), cross-referenced from Phases 3 and 7 and not recomputed here.
- **The machine-readable property taxonomy for CBOMs** is defined by the Applied Quantum CBOM Profile (CBOMProfile.org), cited by minimum version in Phase 2 and not required.
- **Sector-specific enumerations** are carried by the six sector extensions, which follow this framework's sector taxonomy.
- **The regulatory roster** in the Regulatory Timeline Context is a dated snapshot, and this document states the mechanism and three anchor instruments.

The current roster, with the jurisdictions and sector supervisors the table does not cover, is maintained on the Global PQC Migration Clock page and updated through the PQC Migration Brief (Keeping Current, below).

The methodology 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 the phase starts

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

- **Activities** – what the phase does, with enough specificity to assign to a team next Monday
- **Decision Points** – inline callouts where choices branch on context
- **Applying This Phase** – who runs it, the minimum viable version for a small estate, the decision it produces and who takes it, and what it hands to the next phase
- **Outputs** – what the phase produces, with quality criteria
- **Interdependencies** – what feeds into and out of the phase
- **Common Failures** – what goes wrong and how to avoid it, closing with the single failure the phase's outputs most often conceal
- **Challenge Questions** – what the SteerCo asks before accepting the phase output. The assurance questions in GRC Implementation apply the same tests
- **Maturity Indicators** – how to assess whether the phase is done well enough to proceed

Appendix I shows worked artifacts across the phases and at closure, invented and labeled as such, so that a reader can see what each phase's output looks like on paper.

Phases are not executed sequentially in a clean waterfall. Real programs overlap: Phase 1 discovery runs on some systems while Phase 5 pilots run on others. The phases are ordered by dependency. The calendar comes from the Phase 4 roadmap.

NORMATIVE CONVENTION

In normative provisions, **must** denotes a requirement, **should** a recommendation from which a documented departure is permitted, and **may** an option. Examples, sample statements, illustrative thresholds, the reference program economics in Activity 0.2c and the worked artifacts in the appendices are informative. Where this framework takes a position, the position is stated as a *should* with its reason.

Four provisions are stated as *must* because departure defeats the method:

1. The deployment environment classification (Phase 5)
2. The signed accepted-gaps register at gate G1 → G2 (Activity 4.6)
3. The agility criterion at every migration gate (Activity 4.6)
4. The closure proof (Migration Verification & Program Closure)

Activity text also uses *must* where a step is mandatory within its activity. Those requirements bind inside the activity that states them, and the four provisions above are the program-level ones. Appendix J lists every *must* statement in the document, the four program-level provisions first, generated from the text at each build so that the list cannot drift from it.

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

THREE LAYER SYSTEMS

Throughout this document, "the estate" means the organization's cryptographic estate: every system, service, device and data store that uses or depends on cryptography, wherever it runs and whoever operates it.

This document uses three numbered layer systems. Do not read them as one system:

- Discovery Layers 1 to 5 (Activity 1.2).
- CBOM Layers 1 to 4 (the Minimum Viable CBOM model in Phase 2).
- The Cryptographic Concentration Framework's layers 1 to 6, used here by cross-reference only.

Every CCF layer reference in this document includes the prefix and the layer name, as in "CCF layer 4, key custody platform".

ENTRY POINTS BY READER

No single reader needs the entire document at once. The table below names what each reader executes and the artifacts that reader owns. The rest is reference material for when the corresponding work begins.

Reader	Read first	Artifacts owned
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; Appendix F for the objections raised in the funding conversation	Program charter; risk appetite statement; board reporting; the closure decision
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 Re-Baseline Protocol (Activity 4.8); Metrics, KPIs &	Roadmap; PMO operating model; KPI dashboard; re-baseline records; BAU handover register

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

Reader	Read first	Artifacts owned
	Reporting	
Security architect and cryptographic engineer	Phases 1, 2, 5 and 6; Appendix A for algorithm parameters; Appendix H for protocol-family coverage; the Crypto-Agility foundation	Discovery outputs and evidence grades; CBOM architecture; pilot designs; validated migration patterns; rehearsal records
Procurement and vendor management	Phase 7 in full, with the contract language in Activity 7.3, the readiness questionnaire in Activity 7.2 and the cloud and SaaS classification in Activity 7.7	Vendor classification; questionnaire responses graded for firmness; contract clauses; the vendor scorecard
GRC and internal audit	Phase 3; the GRC Implementation foundation; the Regulatory & Standards Alignment Map; Migration Verification & Program Closure; Appendix G for mapping this framework onto frameworks the organization has already adopted	Quantum Readiness Assessment; KRI reporting; Regulatory Horizon Report; evidence dossier; audit records
Crypto-agility workstream owner	The Crypto-Agility foundation in full; the workstream structure in Activity 0.3; the agility-debt register in Phase 3; the gate criteria in Activity 4.6; the rehearsal pattern in Activity 5.2; the closure proof in Migration Verification & Program Closure	The mechanisms catalogue as applied to the estate; the rehearsal schedule and records; agility-debt burn-down; the agility KPIs

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

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

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 by SaaS and cloud	Phases 0 and 4 merge into a single charter-plus-plan; one combined inventory/CBOM workstream; SteerCo folds into an existing security forum; one or two pilots	Risk-driven scoping (1.0); a queryable CBOM; questionnaires to the top 10 vendors; hybrid TLS on internet-facing endpoints
Mid-size (~2,000–20,000)	Thousands of certificates; mixed cloud and on-premises; some OT or embedded	Ten workstreams consolidate to four or five, with the Crypto-Agility and Cryptographic Record owners still named where their streams merge; 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.

KEEPING CURRENT

This edition states standards, regulatory anchors and vendor-independent facts as of its publication date. Between editions, the following are published through the PQC Migration Brief at pqcmigrationbrief.com:

- corrections
- standards and regulatory movement
- verify-first items that resolve
- changes to the anchors and archetypes in the Regulatory & Standards Alignment Map

This document's Version History records what changed at each edition. The Brief is an optional notification channel for readers who want those changes as they happen. The framework itself remains free and ungated at PQCFramework.org.

The deadline roster printed in the Regulatory Timeline Context is the snapshot at this edition. The current roster, with the jurisdictions the table does not cover, is maintained on the Global PQC Migration Clock page on PostQuantum.com and updated from Brief issues.

FRAMEWORK ARCHITECTURE AT A GLANCE

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

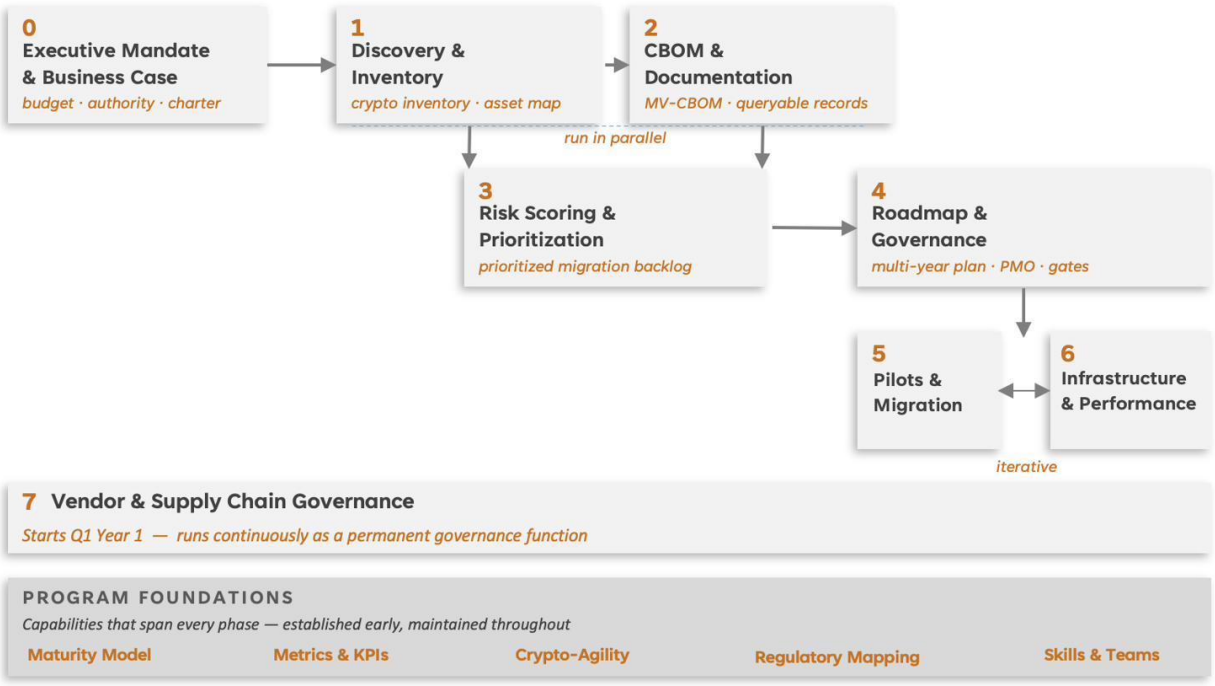

Framework architecture at a glance: eight phases, the parallel vendor governance function and the program foundations

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

Two migration tracks run through Phases 3 to 7: Track A (confidentiality, key establishment) and Track B (integrity and authentication, signatures and PKI). Migration Verification & Program Closure proves the result. Seven program foundations span every phase: the maturity model, metrics and reporting, crypto-agility, regulatory alignment, skills and team structure, SOC implementation and GRC implementation.

REGULATORY TIMELINE CONTEXT

This framework is built against three anchor instruments, one for each deadline archetype the Regulatory & Standards Alignment Map defines:

Archetype	Anchor instrument
Deprecation and disallowance	NIST IR 8547, the reference that other instruments cite
The procurement gate	NSA CNSA 2.0
Plan-by and migrate-by dates set for Member States	The EU coordinated roadmap

The table below is the roster of converging deadlines as of September 2026. Read it with three qualifications:

- It mixes binding requirements, recommended roadmaps and planning milestones at different stages of finalization, so not every date is a legally binding deadline.
- It is a snapshot: instruments and dates move between editions, and the jurisdictions and sector supervisors the table does not cover have dates of their own.
- The current roster, kept up to date and extended to further jurisdictions, is on the Global PQC Migration Clock page on PostQuantum.com, updated from the PQC Migration Brief.

Readers must verify the status and applicability of every instrument for their own jurisdiction and sector before planning against it.

Jurisdiction / Body	Key Milestone	Deadline	Status at this edition
NIST IR 8547	Deprecate quantum-vulnerable public-key algorithms (112-bit security)	After 2030	Initial Public Draft (November 2024); still a draft at this edition
NIST IR 8547	Disallow quantum-vulnerable public-key	After 2035	Initial Public Draft (November 2024);

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

Jurisdiction / Body	Key Milestone	Deadline	Status at this edition
	algorithms		still a draft at this edition
NSA CNSA 2.0 (CNSSP 15, updated 2024 and published March 2025)	New NSS acquisitions CNSA 2.0 compliant, unless otherwise noted	January 1, 2027	Binding for NSS
NSA CNSA 2.0 (CNSSP 15, updated 2024 and published March 2025)	Equipment and services that cannot support CNSA 2.0 phased out, unless otherwise noted	December 31, 2030	Binding for NSS
NSA CNSA 2.0 (CNSSP 15, updated 2024 and published March 2025)	CNSA 2.0 algorithms mandated for use, unless otherwise noted	December 31, 2031	Binding for NSS
NSA CNSA 2.0 (advisory per-category chart, September 2022)	Software/firmware signing uses PQC exclusively	2030	Advisory guidance for NSS. The binding dates are the CNSSP 15 dates above, or the dates the applicable profile notes
NSA CNSA 2.0 (advisory per-category chart, September 2022)	Networking equipment (VPNs, routers) exclusively uses CNSA 2.0	2030	Advisory guidance for NSS
NSA CNSA 2.0 (advisory per-category chart, September 2022)	Web, cloud and operating-system platforms use CNSA 2.0 exclusively	2033	Advisory guidance for NSS. The CNSSP 15 use mandate of December 31, 2031 applies unless a profile notes otherwise
NSA CNSA 2.0 (advisory per-category chart,	Custom applications and legacy equipment updated	2033	Advisory guidance for NSS. The chart carries no date

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

Jurisdiction / Body	Key Milestone	Deadline	Status at this edition
September 2022)	or replaced; niche equipment uses CNSA 2.0 exclusively		beyond 2033, and the CNSSP 15 use mandate of December 31, 2031 applies unless a profile notes otherwise
National Security Memorandum 10 (NSM-10, May 4, 2022)	Mitigate "as much of the quantum risk as is feasible" across federal cryptographic systems	2035	US policy for federal departments and agencies, national security systems included. An overall mitigation goal, not a per-system deadline and not a CNSA 2.0 chart date. OMB M-26-15 carries the same date as its Full Migration phase (below)
Executive Order 14412 (June 22, 2026)	Federal high-value assets and high-impact systems: PQC key establishment	December 31, 2030	Binding for federal agencies; FAR rulemaking directed for covered contractors
Executive Order 14412 (June 22, 2026)	Federal high-value assets and high-impact systems: PQC digital signatures	December 31, 2031	Binding for federal agencies; FAR rulemaking directed for covered contractors
OMB M-26-15 (June 24, 2026)	Agency PQC Migration Plans submitted to OMB and the Office of the National Cyber Director	October 22, 2026 (120 days from the memorandum's own date)	Binding for federal agencies; does not apply to national security systems (44 U.S.C. §3552(b)(6))
OMB M-26-15 (June	Support for TLS 1.3 or	January 2, 2030	Binding for federal agencies, national

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

Jurisdiction / Body	Key Milestone	Deadline	Status at this edition
24, 2026)	a successor		security systems excluded
OMB M-26-15 (June 24, 2026)	Five-phase schedule: Strategy, Planning and Discovery (2026–2027); Pilots and Early Migration (2027–2028); Prioritized Migration (2028–2030); Signature Migration (2031); Full Migration (2035)	2026–2035	Binding for federal agencies, national security systems excluded; plans align with NIST IR 8547 or a successor
Australia (ASD)	Cease traditional asymmetric cryptography in Australian Government systems	2030	Government guidance
NCSC UK	Discovery and planning complete	2028	National guidance
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	April 2026	Federal government; date elapsed at this

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

Jurisdiction / Body	Key Milestone	Deadline	Status at this edition
	developed		edition
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

Whichever roster applies, an organization starting now has between months and nine years to its earliest binding date, depending on jurisdiction and system criticality (the CNSA 2.0 acquisition gate is January 1, 2027. The EU roadmap's first-steps milestone is the end of 2026). A credibly planned migration for a large enterprise requires 4–15 years of execution, which leaves no slack in the schedule.

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

PHASE 0 – EXECUTIVE MANDATE & BUSINESS CASE

PURPOSE

Phase 0 establishes the governance foundation and secures multi-year funding. It is the single most common failure point. Programs that skip or underinvest in Phase 0 stall within 6–12 months when they encounter 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. The scope is broader: applications, integrations, vendor relationships, and data stores that rely on public-key cryptography are all affected. The mandate must frame the program as a multi-year enterprise transformation. Project-level framing produces project-level funding and authority, which do not cover the scope.

The mandate must also state two outcomes together.

1. **Readiness.** The program exists to meet the organization's dated obligations, and every audience-facing artifact (the funding request, internal compliance checks, audit scope, regulatory and insurer reporting, client due diligence) is framed as quantum-safety readiness against those obligations.
2. **The method and the end state.** The program executes through crypto-agility, with the PQC migration as the first exercise of that capability, and it closes only when the organization can change cryptography on a planned date and owns and funds that capability as a standing function (Migration Verification & Program Closure).

State both outcomes in the mandate. That decides on day one whether the organization is running a PQC project or an agility program. Frame the funding request and all external reporting around readiness. Execute the program through agility.

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

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. Quantum Academy (QuantumAcademy.com) delivers the trainings and certifications for the roles this framework describes.

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. The early technical findings provide concrete evidence for the executive commitment the full business case requires.

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 capability, the

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

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.

ACTIVITIES

0.1 Frame the Business Case

Q-Day predictions are speculative, and the business case does not depend on them. Structure the executive argument instead around four "urgency drivers," business pressures that are concrete and current:

- **Regulatory and compliance deadlines.** Map the organization's specific regulatory exposure to the deadline archetypes and enforcement tiers in the Regulatory & Standards Alignment Map, using the table in the Regulatory Timeline Context and, for the jurisdictions it does not cover, the Global PQC Migration Clock page. Three patterns recur:
  - Sector standards already require a documented and annually reviewed record of the cryptographic cipher suites and protocols in use (the payment card industry standard is the common example), which is a cryptographic inventory mandate in all but name.
  - Supervisors in the EU and elsewhere have moved from general risk-management expectations to dated plan-by expectations.
  - For government suppliers, CNSA 2.0 timelines apply through procurement requirements.
- **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 channels, a collector positioned on the path can record the ciphertext and hold it for future decryption. This framework treats HNDL as an active planning risk on that inference rather than on observed harvesting, because a passive collector leaves no evidence in the victim's network. The basis is stated in the Two-Track Migration Model (Phase 5).
- **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.
- **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

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

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 urgent systems are different ones: the trust anchors and signing keys.

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:

- **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 requirement common to all of them.
- **Align with innovation, R&D, and digitization budgets.** Quantum readiness has significant overlap with other strategic investment categories.
    - **Quantum exploration budgets.** 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

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

quantum computing for business applications also need to understand quantum risk.

- **Digital transformation budgets.** 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

Six further arguments strengthen the business case:

- **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
  - misconfigured protocols allowing downgrade attacks

These are real, present-day vulnerabilities that can be remediated immediately. Cryptographic inventories routinely uncover classically vulnerable configurations that

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

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 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.
- **The agility dividend.** Every algorithm change after the first costs less, in the direction Marin's Law states as a design heuristic ($Y \approx K / A$, Crypto-Agility foundation): the readiness budget funds the abstraction, the configuration-driven selection and the rehearsal practice that turn the next mandate into a configuration change rather than a second program. The migration does not consume this part of the spend: it remains after closure as the capability the next transition exercises. The budget request already contains the cost of building agility, and this argument states what the organization keeps once the migration is done.
- **Shared machinery with the AI-security mandate.** Supervisors have begun pairing AI-security and quantum-readiness expectations in the same instruments. The two obligations share their expensive components: the asset and dependency inventory, vendor assurance and contractual controls, and the remediation of legacy systems that cannot be patched. The AI-risk action plan should be built with cryptographic assets in its scope, and the shared work funded once under the readiness heading. The Financial Services Extension contains the banking version of this argument.

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

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

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 (Category 1 in Activity 1.2; vendor names are listed in the front-matter Tool and Vendor References), 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, and 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 (10 streams):

- Discovery (a project with an end: risk-driven scoping, layered discovery, the coverage denominator and the accepted-gaps register)
- Cryptographic Record (the permanent CBOM capability, with a named permanent owner from day one)
- Crypto-Agility (an owner and a budget line from Year 1; scope is the mechanisms catalogue in the Crypto-Agility foundation, the agility-debt register from Phase 3 and the rehearsal program in Activity 5.2)
- Network & TLS/VPN (hybrid rollouts)
- PKI & Code Signing (roots, issuers, toolchains)
- Applications & Platforms (libraries, service mesh, cloud)
- Embedded/IoT/OT (gateways, compensating controls)
- Policy/Compliance/Procurement (standards, clauses)
- Vendor Orchestration (roadmaps, SLAs)
- Education & Change Management (training, comms)

Two changes from earlier editions, each for one reason. Discovery and the Cryptographic Record were one workstream. They are separated because discovery ends and the record does not, and a single workstream responsible for both is

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

disbanded when the first ends, and the second ends with it. Crypto-agility used to be a design principle that every workstream applied and no one owned. It is now a workstream of its own. A capability funded as a principle produces architecture guidance and no owner to apply it. The Crypto-Agility stream has an owner, a budget line and a measured output (Metrics, KPIs & Reporting) from Year 1.

Decision cadence:

- Weekly PMO: track milestones, blockers, vendor responses
- Monthly SteerCo: approve roadmap changes, funding asks, target dates, risk acceptance; ratify the Phase 3 scoring model (dimensions, weights and the appeals rule) before its first run (Activity 3.1)
- Quarterly Board/Risk Committee: KPIs, exceptions, budget status

Decision Point – dedicated program or embedded governance. The program must start as a dedicated program with a named QRPM and a steering function that takes the SteerCo's decisions. An existing operational forum will not give the migration priority at the start. At that point the migration has no inventory, no scored backlog and no vendor commitments.

In the mid-size profile and above, the steering function is a dedicated SteerCo.

A small estate (Right-Sizing the Framework) may run it more simply. An existing security or risk forum may take the role. Roles may also combine: the QRPM may lead a workstream as well, and one executive may both sponsor the program and chair the forum.

One condition applies either way. The forum's minutes name the decisions this framework assigns to the SteerCo as decisions of the program, and name the person who took each:

- roadmap changes
- funding asks
- target dates
- risk acceptance
- ratification of the scoring model and the appeals rule
- gate decisions

A smaller profile compresses the machinery. It never removes the named authority.

The trigger for embedding is maturity Level 4 in the Governance & Ownership and Vendor & Supply Chain domains, or the closure criteria in Migration Verification & Program Closure, whichever comes first. At that point the QRPM role transitions to the standing cryptographic governance owner and the program's continuing capabilities transfer under the BAU handover register. The handover register is the artifact that records the decision and its date.

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

Program Leadership Under the CISO, with the CIO as Co-Sponsor

The most common governance question is whether the PQC migration program should report to the CISO, the CIO, or a joint structure. It 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.

Minimum board engagement model:

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

- **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.
- **CFO pre-brief.** The CFO must be briefed before any board committee sees the funding request, because a multi-year ask that the CFO first meets in the committee papers is contested there. The gross-versus-net split in Activity 0.2c (planned refresh redirected to quantum-safe equipment against net-new money) is the instrument for that conversation.
- **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. The tolerances for HNDL exposure, TNFL exposure, regulatory compliance and crypto-agility are drafted with illustrative thresholds in GRC Implementation (Risk Appetite and Tolerance). The board approves the statement in

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

Phase 0 alongside the charter and the budget, and reviews it annually. The requirement is that the tolerances exist and are specific enough to drive decisions.

Three Lines of Defense

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

First line (the program and the technology teams) owns the cryptographic estate and delivers the migration. The program itself, under the CISO or the QRPM, is first-line management: it plans, executes, measures its own progress and proposes risk acceptances. Application teams migrate their code, infrastructure teams upgrade PKI, HSMs and network devices, procurement updates contract clauses, and each workstream lead is a first-line function.

Second line (the risk function) challenges what the first line delivers: it validates the scoring model and its runs under the model-risk treatment in GRC Implementation, challenges gate decisions and the closure decision, drafts the risk appetite statement for the board's approval, monitors the KRIs, runs the regulatory intelligence function and third-party quantum readiness assessment, and escalates when delivery falls behind plan. A CISO who directs the program is acting as first-line management for that program, whatever the title. The challenge comes from a function that does not deliver the migration. The board approves the risk appetite and receives the KPI pack.

Third line (internal audit) provides independent assurance that the program is governed as chartered and the cryptographic inventory is accurate. It also verifies that the KPIs are measured honestly and that 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.

Cascading Key Risk Indicators

Risk indicators for PQC migration should cascade across three organizational levels. Each level has a different audience, reporting frequency, and depth 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.

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

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 exfiltration detection weighted by confidentiality horizon. 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 kept 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.

0.4 Draft the Program Charter

A one-page charter document covering:

- **Purpose:** Quantum-safety readiness against the organization's dated obligations, delivered through crypto-agility: the PQC migration as the first exercise of a standing capability to change cryptography on a planned date
- **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 applicable regulatory deadlines, with the buffer the risk appetite statement sets.
  - Embed crypto-agility as an organizational capability, proven by a rehearsed algorithm change on Tier-1 systems inside the agility window. The agility

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

window is the 48-hour assessment plus the change window the organization defines, ten business days by default (Crypto-Agility foundation).

- **Discovery access:** Pre-authorization for the Discovery workstream to place passive taps at network choke points, export device and cloud configurations and run credentialed scans across business units, with a sponsor-level escalation path for refusals (Activity 1.4)
- **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:

- Identify the top 20 revenue-generating or mission-critical systems
- For each, determine: primary cryptographic protocols in use, data sensitivity classification, number of dependent systems, vendor ownership vs. internal control
- Estimate the approximate size of the cryptographic estate (number of TLS endpoints, certificates, VPN tunnels, HSM-protected keys, code-signing pipelines)
- 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.

APPLYING THIS PHASE

- **Who runs it.** The Executive Sponsor owns the mandate and the QRPM runs the phase from the day of appointment. The QRPM, with the Policy/Compliance/Procurement workstream and the PMO, produces the charter, the RACI and the board materials, with the CFO briefed before any committee sees the funding request (Activity 0.3).
- **Minimum viable version for a small estate.** Phases 0 and 4 merge into a single charter-plus-plan: a one-page charter in the dual-outcome form (Activity 0.4; Appendix I), a three-year budget line with the umbrella framing of Activity 0.2c, the SteerCo folded into an existing security forum, and the top-20 scoping list from Activity 0.5. The reference program economics are informative. The small estate prices its own list.
- **The decision this phase produces.** The charter approval, the multi-year budget commitment and the QRPM appointment, taken by the Executive Sponsor at gate G0 → G1 and minuted at the governance level that will close the program.
- **What it hands on.** Phase 0 hands four things to later work:
  - The charter, the RACI and the workstream structure, to every phase.

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

- The top-20 scoping list and the top-10 vendor list, to Phases 1 and 7.
- The risk-register entry and the risk appetite statement, to GRC Implementation.
- The board briefing record, to the evidence dossier, as its first entry.

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
Discovery access authorization	Tap placement, configuration export and credentialed-scan authority written into the charter; refusal escalation path named

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 charter (Activity 0.4) provides the authority needed to gain access to systems, repositories, and network infrastructure.

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

- **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 by providing concrete evidence of estate complexity and risk exposure.

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 affects applications, integrations, and vendor relationships across business units. 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 most common 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.

Common failure: a mandate that funds a PQC project and says nothing about the agility capability the program is chartered to leave behind, so that closure has nothing to prove.

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

CHALLENGE QUESTIONS

- Does the charter state both outcomes, the migration and the agility capability, with a success criterion for each?
- Which regulatory archetype and enforcement tier is the business case built on, and which anchor instrument was checked at its primary source?
- Is the budget committed for the multi-year term, or approved for Year 1 with the rest to be requested?
- Who owns the Cryptographic Record after the program closes, and is that person named in the charter today?
- Which of the top-20 systems came from the risk-driven scoping rules of Activity 1.0, and which from a list someone already had?

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 – Progressing	Multi-year budget committed; SteerCo operational; scoping assessment complete
Level 4 – Advanced	Program integrated into enterprise risk register; quantum risk reported to board quarterly alongside other strategic risks
Level 5 – Optimized	Quantum readiness is part of BAU enterprise governance. The QRPM role has transitioned to the standing cryptographic governance owner; risk appetite reviewed annually as a matter of course

© 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. Every subsequent phase depends on this inventory. Prioritization, risk scoring, roadmap construction, and migration sequencing all require knowing which systems use which cryptography.

Major organizations should plan for this phase to take 12–24 months of dedicated team effort. Tool vendors may imply that their products provide a complete solution. No single tool provides 100% coverage. Full inventory requires automated tools, manual audits, and continuous monitoring in combination.

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

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

indefinitely. Organizations should plan staffing accordingly: the discovery team is not a temporary project team, it is a permanent function.

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. When access politics block discovery, coverage gaps remain. Those gaps grow larger through every later phase.
- **At least a preliminary asset inventory or architecture documentation**. Perfect asset data is not required (most organizations could never start if it were), 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.

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

- **Cooperation agreements with OT/ICS teams (if applicable)**. OT environments require their own discovery approaches, tools, and 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.

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)

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

is sufficient to begin Phase 2 CBOM population and Phase 3 risk scoring on the most critical systems, while the broader inventory catches up.

1.1 Establish Three Parallel Inventories

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

- **The 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.
- **The Sensitive Data Inventory (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.
- **The 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 inventories must be coordinated through a single governance structure, even if executed by different teams using different tools. The intersection of all three ("this system (Systems and Assets Inventory) protects this sensitive data (Sensitive Data Inventory) using this vulnerable algorithm (Cryptographic Inventory)") is what enables defensible prioritization in Phase 3. They are inventories and not tracks: Track A and Track B name the two migration tracks of the Two-Track Migration Model (Phase 5) everywhere in this document.

Assigning confidentiality horizons (Sensitive Data Inventory 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 secrecy requirements**, including the data-protection and professional-secrecy duties that run for as long as the data is held (classified material per national rules; medical, legal-professional and banking confidentiality; personal data under data-protection law).
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.
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.

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

The horizon is one of five clocks, and the register records each separately with its owner and its source, because they measure different things and end on different dates:

Clock	What it measures	Owner
Confidentiality need	How long the data has to stay secret: the horizon above	The information owner
Retention obligation	How long the record has to be kept and, in most regimes, kept readable: health records 10 to 30 years depending on jurisdiction, audit and tax records 7 to 10	Records management
Legal preservation period	A litigation hold or a regulatory freeze that suspends deletion	Legal
Signing authority's period of authority	For signed material: the certificate validity or the mandate under which the signer may sign	The PKI or signing-service owner
Verifier's service life	For signed material: how long a relying party will still verify the signature (a fielded device for its service life, an archived contract for as long as it is relied on)	The product or archive owner

A record can be retained for years after its secrecy has expired, and a signature's mathematical verifiability, its signer's period of authority and its verifier's service life are three different durations. Which clock each part of the framework reads:

- **Phase 3 Dimension 1** reads the confidentiality need and nothing else.
- **The data-at-rest strategy (Activity 5.6)** is decided by the retention obligation and the preservation period, which never lengthen the HNDL horizon.
- **Dimension 2's verification-period factor and the Track B timing (Activity 3.3)** read the verifier's service life, or the signing authority's period of authority where every relying party can be updated when it ends, whichever is longer for the anchor. Appendix C reads Mosca's X the same way.

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

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:

Layer 1 – Network Traffic Analysis (Passive) Deploy passive monitoring on network taps or SPAN ports to capture TLS/SSH/IPsec handshakes and extract negotiated key-exchange groups, cipher suites, key sizes and protocol versions, and certificate chains where the protocol still sends them in the clear (TLS 1.2 and earlier; TLS 1.3 encrypts them, so certificate evidence comes from endpoint telemetry and the certificate inventory). 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

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

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 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 capability requirements for the category are: at least three discovery methods in one platform (network, code and configuration); output in CycloneDX CBOM format that includes the minimum record in Activity 2.1; an evidence grade on every finding; and integration with the CBOM platform selected in Phase 2.

Vendor and product names for the category are listed in the front-matter Tool and Vendor References with an as-of date, and on PQCFramework.org. The category changes 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-

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

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 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

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

that produce output in CycloneDX CBOM format or that integrate with the CBOM tooling selected in Phase 2.

Evidence grading. Every discovery finding must carry an evidence grade, in this framework's five tiers:

Grade	Meaning	Example
E1	Runtime observed	A captured handshake, a cipher suite negotiated in production traffic
E2	Binary confirmed	A library identified in the deployed binary
E3	Inventory declared	A CMDB record, a certificate-manager export, an SBOM component entry
E4	Vendor attested	A supplier's answer to a questionnaire
E5	Inferred	Reasoning from documentation or a certification record

The line between the first two tiers is whether the operation was seen executing: E1 is the operation observed at runtime, whether captured on the wire by passive monitoring or observed on the host by runtime instrumentation of the running process; E2 is presence confirmed in the deployed artifact without observing the operation, by binary analysis of the installed executable or firmware or by memory analysis of a loaded image.

Applied to the layers above:

- Layer 1 passive monitoring produces E1.
- Layer 4 produces E1 where runtime instrumentation observes the call executed, and E2 where binary or memory analysis confirms that the library or code path is present.
- CMDB, configuration and certificate-manager exports (Layer 3) produce E3.
- Vendor documentation and questionnaire answers produce E4.
- Every finding, enrichment or classification produced by Category 9 tooling is E5 until a person or a higher-tier method verifies it upward.

Record the grade with the finding. Keep it when the finding is written into the CBOM and when it is scored in Phase 3.

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

A low grade does not stop funding for the work the grade itself calls for. An E5 or E4 entry scored Critical justifies a funded investigation to raise the grade. It also justifies any reversible mitigation the exposure warrants. What it does not justify is accepting the migration as complete. That step needs E1 evidence, or the function's positive evidence in the verification matrix, before the entry counts as migrated (Migration Verification & Program Closure).

The tiers align one to one with the evidence tiers of the Applied Quantum CBOM Profile, so no translation is needed for organizations already using it. Grade each claim, not each inventory row. A claim is one observed fact: an algorithm, a key size, a protocol version, a certificate. A single row usually records several claims, so it usually has several grades. Keep the Phase 3 decision fields and the evidence fields as linked records. Do not collapse them into one graded line.

The coverage denominator. Coverage must be measured against an estate enumerated before discovery ran, never against what the tools found. The denominator is built from the CMDB, the service map and the application portfolio (Activities 1.4 and 1.5), and every coverage figure states it.

- "94 percent of the 2,140 TLS endpoints enumerated from the load-balancer and service registries" is a coverage claim.
- "94 percent coverage" without a denominator is a description of one tool's output.

The Inventory Completeness KRI (GRC Implementation) cannot be measured without the denominator. The denominator will grow as discovery matures. Report that growth as a denominator change and state its cause (Metrics, KPIs & Reporting). Never report it as a coverage regression.

Some asset classes cannot be observed by the discovery method in use. Name those classes as unobserved. Their coverage is unknown, not zero.

Earlier editions asserted that no tool alone achieves completeness, and cited no source for it. The source is the ACM TOSEM study of January 2026 (DOI 10.1145/3788692). The study compared 55,444 SBOMs generated by six tools across 3,287 repositories. Inter-tool package-detection consistency was 7.84 to 12.77 percent, depending on the language. A proof built on one tool's output is therefore only a proof about that tool's output.

Discovery will not reach some of the scope inside the denominator. Record that scope in the accepted-gaps register (Outputs, below). Each entry states the reason and a planned closure date. The register is signed.

Two sibling instruments record related facts. This framework cites them and does not repeat them.

- The Applied Quantum CBOM Profile records coverage per service as `conc:discoveryCoverage`, with its `conc:discoveryMethod`.

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

- The Cryptographic Concentration Framework publishes a Discovery Coverage Attestation at CCFramework.org. It is a signed statement about a bill of materials, not content inside one.

1.3 Map the Cryptographic Estate

For every cryptographic instance discovered, record:

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; weakened by Grover; not known to be quantum-vulnerable; unknown (a legitimate value, never defaulted)	Shor (RSA key exchange), not known to be quantum-vulnerable (AES-256), unknown (undocumented vendor library)
Owner	Responsible team/individual	Platform Engineering / J. Smith

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

Field	Description	Example
Vendor Dependency	Whether migration requires vendor action	Yes – Vendor X controls firmware
Control Posture	Whether organization controls both endpoints	Full control / Partial / No control

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:

- **Configuration Management Database (CMDB).** The starting point only: CMDBs are often stale, incomplete, and biased toward officially provisioned infrastructure. Use as a baseline, then validate and extend.
- **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.
- **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.

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

- **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.
- **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.
- **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.
- **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.
- **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.

- **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.
- **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.
- **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

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

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 quantum readiness program is often what finally forces the asset inventory to completion, producing immediate security value from the inventory work itself.

The facilities estate. Every enterprise operates an OT estate inside its own buildings: access control, building management, elevators, fire and life-safety panels, CCTV and the controllers behind them, almost none of it owned by the security team or present in the CMDB. In one corporate head office, a walkdown identified more than 160 distinct categories of connected device, most of them provisioned by facilities management or the building owner.

That count comes from engagement experience, not from a study, and the population differs from building to building. The general point still applies. Treat the facilities estate as a discovery scope of its own. Engage its owners before the walkthrough (the physical walkthroughs above). Expect its own gateway and encapsulation patterns (OT & Critical Infrastructure Extension).

Access governance. The charter (Activity 0.4) pre-authorizes the Discovery workstream to place passive taps at network choke points, export device and cloud configurations and run credentialed scans, with a sponsor-level escalation path for refusals. Escalate the first refusal within five business days. Do not wait for the next SteerCo. Every week of blocked access adds a coverage gap, and those gaps grow through Phases 2 and 3.

Two calibration observations from engagement experience set expectations for what the access buys:

- Physical walkdowns routinely find 20 to 30 percent of register entries outdated (moved, decommissioned but still running, or replaced by a device the register does not know).
- Passive monitoring at choke points finds more live TLS-terminating systems than the CMDB admits.

Both are reasons to treat the register as the starting hypothesis and observed traffic as the record. Reaching an asset and knowing who owns it are two different problems. The OT & Critical Infrastructure Extension's cryptographic ownership register covers ownership.

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

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
- **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)

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

- **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, with every new application deployment, library update, certificate issuance and rotation, vendor firmware release, cloud configuration change, and new instance of shadow IT. 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 Program Foundations). Crypto-agility requires current visibility into cryptographic posture: without it, an organization 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 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

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

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:

- 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 hooks, 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.

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

- **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 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	Immediate triage; assess whether this is active exploitation risk; remediate or isolate within 72	Same day

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

Alert Level	Trigger	Response	Timeline
	internet-facing TLS endpoint detected	hours	
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
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.

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

- 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
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

APPLYING THIS PHASE

- **Who runs it.** The Discovery workstream runs the layered discovery. It is a project with an end. The Cryptographic Record owner receives what it finds from the first week. The QRPM is accountable for the coverage denominator and the accepted-gaps register. The Discovery workstream lead signs the register, the CISO countersigns it, and the SteerCo accepts it at gate G1 → G2.
- **Minimum viable version for a small estate.** Run one combined inventory-and-CBOM workstream. Apply risk-driven scoping (Activity 1.0) to the top-20 list. Run Layer 1 and Layer 3 discovery (network-negotiation and certificate data, Activity 1.2) on Priority A systems inside 60 days. Limit Layer 2 code scanning to the in-house

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

applications that handle long-horizon data. Build the denominator from the asset register and the cloud account list, not from a discovery tool's own count.

- **The decision this phase produces.** The signed accepted-gaps register: which parts of the estate are not yet inventoried, why, and who owns closing each gap by when, taken by the SteerCo at gate G1 → G2. Do not pass the gate on coverage alone. A coverage figure describes what the tools found, not what the estate contains.
- **What it hands on.** Phase 1 hands four things to later work:
  - Discovery findings with their evidence grades, to Phase 2 for the minimum record.
  - The coverage denominator and the accepted-gaps register, to Phase 4 (gate G1 → G2) and to the Metrics foundation (Inventory KPI).
  - The vendor dependency data, to Phase 7.
  - The continuous-discovery infrastructure, to SOC Implementation and to Migration Verification. Its observations become the evidence standard there.

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 horizon, retention obligation and any preservation period, recorded as separate clocks with their owners (Activity 1.1)
Systems and assets register (integrated)	All assets classified by criticality and vendor dependency; cross-referenced with BIA classifications
Accepted-gaps register (signed)	Scope inside the coverage denominator that discovery did not reach, listed with reason,

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

Output	Quality Criteria
	owner and planned closure date; signed by the Discovery workstream lead, countersigned by the CISO and accepted by the SteerCo; a criterion at gate G1 → G2 (Activity 4.6)
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

Report classical findings on their own line. Discovery will find deprecated algorithms, weak keys, expired certificates and downgrade-capable configurations. Route their remediation to security operations, as you would today, and report that work as a parallel benefit stream. Never count it toward PQC migration KPIs. A backlog that reports classical remediation as PQC progress overstates readiness by the size of the easiest work.

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).

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

- **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.

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

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

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.

- **Disbanding discovery after the initial inventory.** The initial inventory is the project with an end; continuous discovery is the standing capability that follows it. Organizations that disband the whole team after "finishing" the initial inventory find their CBOM becomes stale within 6 months.

Common failure: an inventory that cannot answer a question this afternoon: which systems negotiate RSA key exchange on an internet-facing endpoint, and which of them carry data with a confidentiality horizon past 2035.

CHALLENGE QUESTIONS

- What is the coverage denominator, how was it built, and what changed it since the last report?
- Which findings carry evidence grade E1 or E2, and which are inferred from configuration or vendor documentation alone?
- Which gaps are in the signed register, who owns each, and which have passed their closure date?
- Which Priority A systems were inventoried first, and by what rule were they chosen?
- What did continuous discovery detect this month that the initial inventory missed?

MATURITY INDICATORS

Level	Indicator
Level 0 – Unaware	No cryptographic inventory exists; no awareness of the need
Level 1 – Aware	Partial manual inventory of obvious systems (web servers, VPN); CMDB-only asset register; no continuous discovery
Level 2 – Initiated	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 – Progressing	≥90% coverage; continuous discovery in CI/CD and passive monitoring; integrated with

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

Level	Indicator
	CMDB, SBOM, BIA, and certificate management; change management integration live; crypto champions designated; alerting framework operational
Level 4 – Advanced	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
Level 5 – Optimized	Discovery runs as a funded standing capability under the Cryptographic Record owner; coverage measured against a maintained denominator; the accepted-gaps register closed or re-accepted on schedule

© 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. Risk scoring (Phase 3), roadmap construction (Phase 4), and migration verification all operate against the CBOM.

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 asked in every subsequent phase: 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 1.7 is the format (see Activity 2.1; 1.6 is accepted during a documented transition), but the organization must also decide where the CBOM will be kept: 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.

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:

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

- **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).

ACTIVITIES

2.1 Select CBOM Format and Tooling

Format. The CBOM uses CycloneDX v1.7 (October 2025), ratified as ECMA-424 second edition (December 2025), which defines cryptographic asset types for algorithms, certificates, keys, protocols and related cryptographic material, and the `dependencies` relationships between them. CycloneDX 1.7 also introduced the Cryptography Registry, created because inconsistent algorithm naming across tools was blocking PQC readiness assessment in deployed CBOMs.

A CBOM that records the registry identifier, or the algorithm OID where one exists, is countable across tools. One that records tool-specific names is not.

Three 1.7 version notes govern producers and consumers:

- `ellipticCurve` replaces the deprecated `curve`.
- `relatedCryptographicAssets` replaces the deprecated `certificateProperties.signatureAlgorithmRef`.

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

- `securedBy` is present and is not deprecated. A document produced under 1.6 is accepted during a documented tooling transition, with `curve` and `signatureAlgorithmRef` accepted for compatibility and the move to 1.7 planned with the tooling refresh.

The Linux Foundation PQCA's CBOMkit is the open-source generation and viewing kit at this edition. Commercial CBOM tooling is listed in the front-matter Tool and Vendor References.

The record this framework specifies is complete on its own. The table below defines, in this framework's own terms, what a CBOM entry must record, and a CBOM built to it can be expressed as plain CycloneDX. Organizations that adopt the Applied Quantum CBOM Profile (CBOMProfile.org), v1.0-RC or later, carry the same record as Profile properties under the `appliedquantum` CycloneDX property namespace. The Profile is the machine-readable carriage of what this framework asks for. It is useful and it is not required, and a reader who never opens it can execute every phase.

Whichever carriage is chosen, the Minimum Viable CBOM architecture (the four CBOM layers above, freshness governance in Activity 2.4, queryability in Activity 2.3) is defined by this framework. The namespace registration is CycloneDX property-taxonomy issue #189, filed August 14, 2026, reserved and not yet in the published taxonomy table. It is cited with that status and is not described as registered until the registry row is published.

The `pqc:`, `custody:`, `entropy:` and `conc:` prefixes used in this document are the Profile's property groups, written as shorthand for their fully qualified names under the reserved namespace. The Profile specification gives the full names.

CBOM minimum record (this framework's definition):

A CBOM built to this table records the information the Profile's Tiers 1 and 2 require. An organization that adopts the Profile claims conformance to it only after validation against the pinned Profile version and its rules. An organization that does not adopt it owes no conformance claim.

Field	Purpose	Informative: CycloneDX / Profile carriage
Component identifier	Links to asset inventory and CMDB	Native bom-ref; `pqc:cmdbRef`
Algorithm OID	Unambiguous algorithm identification	`cryptoProperties.oid`, with the Cryptography Registry identifier where one exists
Key size / security parameter	Determines quantum vulnerability class	`algorithmProperties.parameterSetIdentifier`, `.classicalSecurityLevel`,

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

Field	Purpose	Informative: CycloneDX / Profile carriage
		`.nistQuantumSecurityLevel`
Protocol context	Where the algorithm is used (TLS, IPsec, S/MIME, etc.)	`protocolProperties.*`
Implementation (library + version)	Identifies patching and upgrade path	Native component, tied to the algorithm through `dependencies.provides`
Certificate reference	Links to certificate chain for PKI-related crypto	`certificateProperties.*`, with the certificate-signature algorithm via `relatedCryptographicAssets`
Cryptographic function	Key establishment / digital signature / symmetric / hash; routes the entry to Track A or Track B (Activity 3.3)	`algorithmProperties.cryptoFunctions`
Confidentiality horizon (years)	The confidentiality need from the Phase 1 classification (Activity 1.1), not the retention obligation; read by Phase 3 Dimension 1	`pqc:dataRetentionPeriod`
Regulatory deadline (date and source instrument)	The date after which the algorithm is disallowed for this entry, and the instrument that sets it; read by Phase 3 Dimension 4	`pqc:disallowedAfter`; `pqc:horizonSource`
Data classification	From the Sensitive Data Inventory (Phase 1, Activity 1.1)	`pqc:dataClassification`
Quantum vulnerability status	Vulnerable (Shor), graded by quantum attack-cost band per algorithm and key	`pqc:vulnerabilityStatus`

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

Field	Purpose	Informative: CycloneDX / Profile carriage
	size (see Phase 3); weakened (Grover); not known to be quantum-vulnerable; unknown. Unknown is a legitimate value and is never defaulted to not-vulnerable	
Migration status	Not started / Planned / In progress / Hybrid deployed / PQC-only / Not applicable	`pqc:migrationStatus`
Owner	Responsible team for migration decisions	`pqc:systemOwner`
Vendor dependency flag	Whether migration is self-controlled or vendor-dependent	`pqc:vendorDependency`
Trust scope (recorded at Phase 3)	The reach of the trust the key or certificate anchors: enterprise-wide root, shared intermediate, single-application signing key; read by Phase 3 Dimension 2	No Profile property records the scope at v1.0-RC; `pki:trustAnchorRef` records the anchor relied upon and `conc:serviceRef` the services that depend on the asset, and the scope is read from them
Verifier updatability (recorded at Phase 3)	Whether the verifiers of this key's signatures can follow an algorithm change: centrally managed trust stores, updatable by campaign, or not updatable; read by Phase 3 Dimension 2 and by the agility-debt register	No Profile property at v1.0-RC

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

Field	Purpose	Informative: CycloneDX / Profile carriage
Dependency firmness (recorded at Phase 3)	The firmness grade of the vendor or counterparty date the migration waits on: committed, forecast, announced, unknown (Activity 7.5); read by Phase 3 Dimension 3	`pqc:vendorReadinessStatus`, with the date in `pqc:vendorReadinessDate`
Compensating controls (recorded at Phase 3)	The interim controls in place for the entry (segmentation, overlay encryption, key-lifetime reduction, a bridging pattern under Activity 7.4) and the exposure they leave; read by Phase 3 Dimension 3 and carried into the verification record's residual-exposure column	`pqc:compensatoryControls`
Delivery state (recorded at Phase 3)	Available, dependent, blocked or unknown (Activity 3.2); read by Activity 4.2 for roadmap placement, by Activity 4.6 gate records and by Activity 7.4 for the bridging decision; never by the tier test	No Profile property at v1.0-RC; `pqc:vendorDependency` and `pqc:vendorReadinessStatus` carry the dependency the state is read from

- **Five fields recorded at Phase 3.** The last five rows are populated at Phase 3, by the scoring run and from the vendor register, and are not Phase 1 discovery findings. The record leaves them empty until then, and a vendor-supplied CBOM (Activity 7.2) is not expected to fill them. They are part of the minimum record because Phase 3 reads every one of them: a CBOM built without the columns cannot be scored on Dimensions 2 and 3, and verifier updatability is the strongest single indicator of agility debt in the estate (Activity 3.1).

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

- **The unit of a row.** A row of the record is one cryptographic use: a component, the operation it performs (the cryptographic function, with its algorithm and parameters), and the context it performs it in (the protocol context or store, and the peer or trust relationship). A component with two protocol contexts is two rows, and a key-establishment operation and the certificate that authenticates it on the same interface are two rows (Appendix I.3).

Endpoint counts and component counts are derived from the rows and never collected separately: an endpoint is the set of rows that share one component identifier and one network-facing protocol context, and a component is the set of rows that share one component identifier, so the Coverage KPI's Tier-1 endpoints and the Phase 3 scored-entry count are read from the same population (Metrics, KPIs & Reporting).

- **Forward alignment with the federal CBOM guidance mandate.** Executive Order 14412 §5(d) directs CISA, coordinating with NIST, to publish CBOM minimum-elements guidance within 270 days of June 22, 2026, so by March 19, 2027 (White House text; Federal Register Vol. 91 No. 121). The direction compels no CBOM production and says nothing yet about content. Two consequences apply. When the guidance publishes, the minimum record above and the Profile citation are reviewed against it and the differences are recorded in the Version History rather than absorbed silently.

Until then, the CBOM has moved from a best practice to an instrument with a named guidance mandate behind it, which is the strongest sentence available to the Phase 2 funding case.

2.2 Populate CBOM from Inventory Data

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

- Import automated discovery results into CycloneDX format using tool-native exporters or CBOMkit
- Enrich with manual investigation findings for systems not covered by automated tools
- Cross-reference with SBOM data to establish library dependency chains
- Link to certificate management system for PKI-related entries
- Add data classification tags from the Sensitive Data Inventory (Activity 1.1)
- Flag vendor dependencies from the Systems and Assets Inventory (Activity 1.1)

2.3 Integrate CBOM into Operational Processes

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

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

- **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
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, which of them make up the unmigrated long tail, which vendor products block remediation, and which data stores have 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.

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

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 be classified at the organization's most restrictive standard level (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, and never standing access. Consultant access ends with the engagement.
- **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.

Mark the document at generation. CycloneDX 1.7's `metadata.distributionConstraints` sets a document-level TLP marking whose default is CLEAR. A CBOM must carry a TLP marking set at generation, chosen by the recipients it may be shared with, and is never left at the default. TLP is a sharing rule and not a classification: it states who may receive the document, while the organizational classification in the first bullet states how the document is protected.

The two are recorded separately. Under FIRST TLP 2.0 the levels narrow from CLEAR through GREEN, AMBER and AMBER+STRICT to RED. AMBER permits need-to-know disclosure inside the recipient organization and with that organization's clients. AMBER+STRICT permits disclosure inside the recipient organization only. RED permits disclosure only to the individuals named when the document is shared, and those individuals may sit at a vendor.

The full CBOM handled under the Restricted classification is marked at the level that matches the recipients the organization intends: RED for named individuals on the program team, AMBER+STRICT for a vendor or auditor extract the recipient may circulate internally but no further, AMBER where the recipient's own clients are in scope. A marking that disagrees with the intended sharing model, in either direction, is a marking error and not extra caution.

Where a contract or an NDA restricts sharing more narrowly than the marking permits, the contractual restriction governs and is recorded with the marking. The furthest any regulator has yet gone on bill-of-materials handling is CERT-In's Technical Guidelines v2.0 (July 9, 2025): encrypted storage and transmission, a single access-controlled

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

inventory, and access logging, which this framework adopts as the handling floor. The vendor-side clause, covering how a supplier handles the CBOMs it receives and produces, is in Activity 7.3.

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

Attest, do not send. When a supervisor, auditor, insurer or counterparty asks for the CBOM, the default response is a signed commitment to the snapshot (a hash and its signing record), a coverage statement that gives its denominator and discovery method (Activity 1.2), and the predicate the requester needs (for example, "no internet-facing Tier-1 endpoint negotiated classical-only key establishment during the observation window"), in place of the document.

The requester receives the answer to its question and a proof that the answer was drawn from a fixed record, and the organization's map of its own weak points stays inside. The signed commitment proves that the record was fixed at a date. It does not prove that the record is complete or that the predicate was evaluated correctly against it, so the attestation states coverage alongside it with its denominator and why the population, the query and the evidence behind each predicate are retained for examination (Signed snapshots, below).

Where the requester has examination rights, the point-in-time extract through a controlled channel under the sharing rule above is the form of access: "attest, do not send" limits what leaves by default and is not a refusal of examination. The Cryptographic Concentration Framework's Discovery Coverage Attestation (CCFramework.org) is the family's published predicate of this kind: a signed statement about a bill of materials, never bill-of-materials content.

Signed snapshots. Where auditors will ask how a snapshot is proven unaltered, RFC 9943 (the SCITT architecture, Standards Track, June 2026) is the standardized form: an append-only transparency service, COSE-signed statements and receipts that commit a CBOM snapshot at a date. The quarterly CBOM snapshots in the evidence dossier (Metrics, KPIs & Reporting) should be committed this way where the capability exists, and hashed and signed under the organization's document-signing standard where it does not.

A signature or a receipt establishes integrity and registration of the snapshot. It does not establish that the inventory is complete or that a figure derived from it was calculated correctly, which requires the retained population, query and evidence behind the figure. The attestation described above therefore states coverage beside the commitment. The published precedent for the accepted-gaps register in Phase 1 is the "Explicitly Identifying Unknown Information" element of CISA's 2026 Minimum Elements for an SBOM (July 29, 2026), which distinguishes unknown from withheld and renames the 2025 "Known Unknowns" element.

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

APPLYING THIS PHASE

- **Who runs it.** The Cryptographic Record workstream populates and governs the CBOM, with the permanent owner named in Phase 0. The Applications & Platforms workstream owns the CI/CD and change-management hooks of Activity 2.3. The QRPM is accountable for freshness governance, and the CISO for the artifact-security controls of Activity 2.5.
- **Minimum viable version for a small estate.** Build a queryable CBOM in CycloneDX that records the minimum record of Activity 2.1 for Priority A systems. Generate it with the open-source kit named in Activity 2.1, or export it from the discovery tooling. Populate Layer 1 and Layer 2 of the Minimum Viable CBOM model first. Take one signed snapshot per quarter. Limit access to the roles named in Activity 2.5. A spreadsheet that cannot be queried by algorithm, protocol and confidentiality horizon is an inventory, not a CBOM.
- **The decision this phase produces.** The CBOM structure, format and freshness rule, approved by the SteerCo at gate G1 → G2, and the declaration at gate G2 → G3 that the CBOM is populated to the depth Phase 3 needs, taken by the QRPM on the Cryptographic Record owner's report.
- **What it hands on.** Phase 2 hands four things to later work:
  - The minimum record per entry, with evidence grades and the three Phase 3 input fields (cryptographic function, confidentiality horizon, regulatory deadline with source), to Phase 3. The five fields marked as recorded at Phase 3 are left empty for Phase 3 to populate.
  - The CBOM, as the live operational record, to Phases 5, 6 and 7.
  - The signed snapshots, to the evidence dossier.
  - The vendor CBOM requirement and the handling terms, to Activity 7.2 and 7.3.

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

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

Output	Quality Criteria
	entries
CBOM governance policy	Freshness rules, ownership, update triggers documented and enforced
CBOM protection controls	Restricted classification applied; TLP marking set at generation; role-based access and query logging operational; external sharing rules documented; snapshot commitment method (signed transparency receipt or signed hash) in use for dossier snapshots

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.

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

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.

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.

Common failure: a CBOM populated to the depth the tools produced rather than the depth Phase 3 needs, so that scoring stalls on entries with no confidentiality horizon and no cryptographic function recorded.

CHALLENGE QUESTIONS

- Which CBOM entries carry all three Phase 3 input fields, and how many entries will be left unscored as unknown exposure?
- When was the last signed snapshot taken, who holds the signing key, and who can read the CBOM today?
- Which changes to production cryptography in the last quarter reached the CBOM through the change-management hook, and which were found by re-scan?
- Which vendor CBOMs received so far conform to the minimum record, and which were accepted in another form?
- Where does the CBOM link to the SBOM for shared libraries, and which library entries have no SBOM context?

MATURITY INDICATORS

Level	Indicator

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

Level	Indicator
Level 0 – Unaware	No CBOM exists; cryptographic documentation is ad hoc or absent
Level 1 – Aware	Partial CBOM in spreadsheet form covering known systems; no standard format
Level 2 – Initiated	CycloneDX CBOM operational for Layers 1–2; queryable; SBOM linkage established for key applications
Level 3 – Progressing	CBOM covers Layers 1–3; integrated into CI/CD; freshness governance enforced; change management integration live
Level 4 – Advanced	CBOM is a real-time operational asset; auto-updated on deployment; Layer 4 gaps systematically managed through vendor governance; CBOM drives automated compliance reporting
Level 5 – Optimized	The CBOM is the operational record of the estate; every production change reaches it through the change hook; signed snapshots and attestations answer external requests without sending the record

© 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. Quantum-vulnerable cryptography varies in risk and in migration difficulty, and 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 migration sequencing follows consistent criteria applied the same way to every entry, with the reason for each decision recorded.

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 whose confidentiality need runs past 10–15 years (medical records, financial archives, government records, legal documents, whose retention obligations usually run at least as long) 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.

The SteerCo must ratify the dimensions, the weights, the ordered tier test of Activity 3.2 and an appeals rule before the first scoring run (Activity 0.3). The appeals rule states who may dispute a tier assignment, on what evidence and to whom, and it is approved with the model rather than after the first dispute. A business owner who contests a tier then appeals under the pre-approved rule. The model itself is not relitigated per system. Weight changes after ratification are minuted under the model-risk treatment in GRC Implementation.

Score each CBOM entry across four dimensions:

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

Factor	Weight	Scoring
Data confidentiality requirement (years)	High	Over 15 years = Critical; over 10 to 15 = High; over 5 to 10 = Medium; 5 or fewer = Low. The years are the remaining confidentiality need recorded under Activity 1.1, never the retention period
Data volume / flow rate	None (sequencing only)	Not a scored factor: it does not enter the highest-of-factors rule and cannot set the dimension's band. It orders entries that tie on the composite score inside a tier (Activity 3.2): the higher-throughput channel (bulk replication, batch file transfer, high-volume API

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

Factor	Weight	Scoring
		traffic, read from the passive-monitoring volumes of Activity 1.6) is sequenced first, because it yields more harvestable material per day of exposure
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 (20 years), or any key with 15 years or more remaining = Critical; intermediate CA (10 years), 5 to under 15 years = High; end-entity (1 year), 1 to under 5 years = Medium; under 1 year (short-lived automated certificates, session keys) = Low
Scope of trust dependency	High	Enterprise-wide PKI root, or a trust anchor relied on across business units = Critical; shared intermediate, or a key whose signatures several applications or external parties rely on = High; single-application signing key = Medium; key whose signatures are verified only by the system that produced them, within one session or process = Low

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

Factor	Weight	Scoring
Signature verification period	Medium	The period over which relying parties will still verify the signature (Activity 1.1: the verifier's service life, or the signer's period of authority where every relying party can be updated when it ends): firmware signed for a 15-year lifecycle, 15 years or more = Critical; 5 to under 15 years = High; 1 to under 5 years = Medium; session authentication, under 1 year = Low
Verifier updatability	High	Verifiers cannot be updated (burned-in trust anchors, unmanaged device fleets) = Critical; updatable by campaign = Medium; centrally managed trust stores = Low

Verifier updatability changes the migration path, not only its urgency. A signing key whose verifiers cannot follow an algorithm change makes key migration insufficient on its own, because the new signatures will not validate until a device campaign or a replacement reaches the verifiers. The OT & Critical Infrastructure Extension treats the non-updatable verifier population in depth; within this phase it is also the strongest single indicator of agility debt (below).

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

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

Factor	Weight	Scoring
		vendor support = Medium; No standard yet = Hard
Interoperability constraints	Medium	Internal-only = Easy; Partner network = Medium; Public internet = Hard (proven at scale; the client population is outside the organization's control)
System age and architecture	Medium	Modern microservices = Easy; Monolithic but maintained = Medium; Legacy unsupported = Hard

Vendor count is not dependency count. Where the feasibility question reaches how many distinct implementations, trust anchors, custody firmware families or key-generation designs execute beneath the vendors an entry depends on, the instruments are the Cryptographic Concentration Framework's Cryptographic Concentration Index per layer per service and its failure-domain reach (CCFramework.org), and this framework does not recompute them. A Tier-1 backlog gated on one implementation family behind four vendors is a CCF finding to be cross-referenced in the QRA, not a Phase 3 score.

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. "Mandatory" means an instrument at enforcement tier 1 to 3 in the Regulatory & Standards Alignment Map; "general expectation" means tier 4 or 5
Audit exposure	Medium	External audit scope = High; Internal audit scope = Medium; Not audited = Low

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

Factor	Weight	Scoring
Contractual requirement	Medium	Customer/partner contract requires PQC = High; Expected soon = Medium; Not required = Low

Scoring inputs from the cryptographic record. Each factor reads named fields of the Activity 2.1 minimum record, or a clock of the Activity 1.1 register that the record links to, so scoring is mechanical once both are populated. The five fields the record marks as recorded at Phase 3 are populated by the Cryptographic Record workstream from the scoring run and the vendor register before the factor that reads them is scored:

Factor	Record fields consumed
Dimension 1 – HNDL risk	Confidentiality horizon; Data classification; Protocol context; the HNDL-exposure determination from Activity 1.1
Dimension 2 – TNFL risk	Certificate reference (validity); Cryptographic function = digital signature; Trust scope; Verifier updatability. The verification-period factor also reads two clocks recorded under Activity 1.1 against the signed material: the verifier's service life, and the signing authority's period of authority. The CBOM record links to those two clocks. It does not store its own copy of them.
Dimension 3 – Migration feasibility	Vendor dependency flag; Dependency firmness (the roadmap firmness grade from Activity 7.5: committed, forecast, announced, unknown); Compensating controls
Dimension 4 – Regulatory pressure	Regulatory deadline (date and source instrument) and its enforcement tier
Track routing	Cryptographic function: key establishment routes the entry to Track A, digital signature to Track B (Activity 3.3)
Roadmap placement and gate records	Delivery state (Activity 3.2), read with the tier and never in place of it (Activities 4.2, 4.6 and

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

Factor	Record fields consumed
	7.4)

Dimension 1's accessibility factor is scored against a state-level collector with owned interception capability rather than an opportunistic attacker: "Internal" means internal to a network such a collector has not already reached, and a flow that crosses a carrier or a cloud boundary scores as partner-facing at least.

Unknown exposure. An entry whose confidentiality horizon or data classification is missing must be left unscored and flagged "classification missing". Never default a missing value to Low, to short-lived or to Public. An entry is therefore in one of two states and never in a third: scored, with every applicable factor read from the record, or unscored, with an incomplete assessment that names the missing field, the classification owner, a review priority and a resolution date.

An unscored entry has no score and no tier. It is not placed on the tiered backlog, and it is counted on its own line in the QRA, never inside a coverage or backlog figure. Two rules apply to unscored entries. The review priority follows the accessibility factor, which is known even when the horizon is not: an unscored entry on an internet-facing or partner-facing system is reviewed first and its owner is asked for the classification inside the current quarter, because exposure is high regardless of what the channel protects. That priority is a queue position and not a provisional score.

No entry may remain unscored past the next quarterly re-scoring: the count and age of unscored entries is reported as a CISO-level KRI (GRC Implementation), and any unscored entry on an internet-facing or partner-facing system, or on a Tier-1 candidate from the Activity 1.0 scoping, that is older than one quarter is escalated to the classification owner.

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), the reference cost published by Roetteler, Naehrig, Svore and Lauter in 2017 (arXiv 1706.06752) put the ECC attack at roughly 2.6 times fewer logical qubits (2,330 against 6,146) and two orders of magnitude fewer Toffoli gates (about 1.3 × 10^11 against about 1.9 × 10^13) than the factoring attack. Even against RSA-2048, a weaker classical target, the 256-bit-curve estimates are in the same band as the factoring estimates.

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

ECDLP attack circuits had received far less optimization attention than RSA factoring circuits, which is why the ECC estimates were the ones more likely to improve. In 2026 they did.

Estimate	Logical qubits for a 256-bit-curve attack	What it rests on
Google Quantum AI, March	1,200 to 1,450 depending on the gate budget	Fewer than 500,000 superconducting qubits
Schrottenloher, June (arXiv 2606.02235)	1,462	Reproduced in the open, about 58 million Toffoli gates for secp256k1
Luo et al., July (arXiv 2607.13816)	835	A circuit that pays for its width in gates
IonQ, submitted September 4, 2026 (Häner, Tripier, Young et al., arXiv:2609.05625, fourteen authors, all at IonQ; the manuscript is dated September 3)	1,457	The secp256k1 attack at about 39 million Toffoli gates consuming 273 million T-states, on a modeled, not built, trapped-ion machine of 19,397 physical qubits (an optimization of IonQ's Walking Cat architecture over qLDPC codes) at 25.7 days per attempt, with a heuristic single-run success probability of 63.3%

The IonQ paper's rigorous claim is a lower bound on logical-level success probability holding with confidence at least 1 − 2^−128, and the 63.3% figure is the heuristic single-run number.

Two limits on reading the table:

- The RSA figures did not move in the same period.
- The IonQ figures are specific to secp256k1, and no estimate at that depth exists yet for P-256 or Ed25519. The assumption ledger behind each 2026 result is analyzed on PostQuantum.com (the IonQ secp256k1 resource estimate, September 8, 2026).

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) have at least as much

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

quantum urgency as RSA-legacy estates. Treat ECC-256 and RSA-2048 as one attack-cost band. RSA-3072 and above is the more expensive target on the published estimates.

None of that assigns a tier. The tier comes from the ordered test of Activity 3.2, which reads dimension scores and factor values from the record. Neither the algorithm nor the key size is among them: an entry protected by RSA-3072 takes the tier its data, its exposure and its trust role give it.

Like the data-volume factor in Dimension 1, the attack-cost evidence is unscored, and it reaches the program in three other ways:

1. **Sequencing within a tier.** It orders entries that tie on the composite score inside a tier, with the larger classical key sequenced later where nothing else separates them.
2. **Read with protocol context and key lifetime.** It is read with the protocol context and the key lifetime recorded for the entry, because a long-lived key that opens every session recorded under it is a different exposure from a per-session ephemeral key of the same size.
3. **Evidence for Mosca's Z.** It is evidence for the Z estimate in Mosca's Inequality (Appendix C), where the algorithm with the lower attack cost is the one that becomes feasible first.

Record the algorithm and key size in each CBOM entry's quantum vulnerability field (Activity 2.1) so all three read one recorded value rather than a judgement re-debated per system.

Agility-debt register. Every Tier-1 and Tier-2 entry scored Hard on Dimension 3's control-posture or system-age factors, or Critical on verifier updatability, is listed in an agility-debt register with the reason: hardcoded algorithm, non-updatable verifier, vendor-locked configuration, or a library with no provider pattern. The register is derived from the scores already recorded and adds no dimension. It feeds the Crypto-Agility workstream (Activity 0.3) and the Phase 4 roadmap, because a system migrated to PQC without removing its agility debt will be migrated a third time at the next algorithm change, and the roadmap must schedule the debt removal with the migration rather than after it.

The register's burn-down is an agility KPI (Metrics, KPIs & Reporting), and before program closure it must reach zero unaccepted debt: every remaining entry funded with a date or under a SteerCo-signed exception (Migration Verification & Program Closure).

3.2 Calculate Priority Scores

For each CBOM entry, calculate a composite priority score. The score is an ordinal sequencing position: it orders entries within a tier and it does not assign the tier (Decision Point, below).

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

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

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.

The arithmetic, stated once. The default mapping is Low = 1, Medium = 2, High = 3, Critical = 4, and Easy = 1, Medium = 2, Hard = 3 for the inverted feasibility dimension. A dimension takes the highest score among its factors: one Critical factor makes the dimension Critical. A dimension that does not apply to an entry (TNFL for a data store with no signature role; HNDL for a signing key that protects no confidentiality) is excluded, and the composite is taken over the applicable dimensions with the ratified weights renormalized to sum to one.

With all four dimensions applicable the composite runs from 1.00 to 3.85, not to 4, because the inverted feasibility dimension caps at Hard = 3 (4 × 0.85 + 3 × 0.15 = 3.85). With Dimension 1 excluded, as for a signing key that protects no confidentiality, the renormalized weights are 0.25 / 0.65 = 38.46 percent for TNFL, 38.46 percent for regulatory pressure and 0.15 / 0.65 = 23.08 percent for feasibility: TNFL and regulatory pressure tie, and TNFL leads the composite only where the regulatory dimension scores lower on that entry, which is a fact about the entry and not about the arithmetic.

An unknown value is not scored: the entry is listed as unscored, the two rules on unknown exposure in Activity 3.1 govern its interim treatment, and it is resolved before the entry is tiered. The SteerCo ratifies the mapping with the weights (Activity 0.3); Appendix I.3 shows the calculation.

Delivery state. Every scored entry records a delivery state with four values: available (the organization can make the change), dependent (the change waits on a delivery with a dated commitment), blocked (no path exists on the roadmap's date), unknown. The state is a separate field from the tier. An entry keeps the tier the ordered test gives it whatever its delivery state. A blocked Tier-1 entry is shown as blocked, with the dependency that unblocks it and that dependency's firmness grade (Activity 7.5), and it appears on the roadmap against its unblocking date (Activity 4.2). It is never demoted for being blocked.

Decision Point – Tier Assignment:

The tier is assigned by the ordered test below, read from the top: the first condition that matches assigns the tier, and the conditions that follow it are not consulted. Every condition is a dimension score or a factor value already in the record (Activity 3.1), so two assessors applying the test to the same record reach the same tier. The composite score plays no part in the test. Dimension 3 and the delivery state play no part in it either: a hard or blocked migration keeps the tier its risk gives it.

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

Tier	Condition (the first that matches assigns the tier)	Target Timeline
Tier 1 – Immediate	Dimension 1 Critical with accessibility High or Critical (an internet-facing entry, or a partner-facing entry with a confidentiality horizon over 15 years); or Dimension 2 Critical on any factor (a root, an enterprise-wide trust anchor, a verification period of 15 years or more, or verifiers that cannot be updated); or Dimension 4 Critical on the specific-regulatory-deadline factor (a mandatory deadline within 2 years)	Begin within 6 months; verified protection in the approved target state for the function within 12 months (hybrid key establishment on Track A; the approved post-quantum signature on Track B, standalone where the profile requires it), which is the framework's default planning target where no external obligation sets an earlier date. An entry that cannot reach it inside the window records its disposition: migration on a dated commitment, enclosure or compensation under an accepted exception with the bridging pattern of Activity 7.4 until the dependency lands, or documented residual risk
Tier 2 – High Priority	Dimension 1 Critical or High, unless the entry is air-gapped (accessibility Low); or Dimension 2 High; or Dimension 4 High on the specific-regulatory-deadline factor (a mandatory deadline within 5 years) or on the contractual-requirement factor (a customer or partner contract that requires PQC)	Begin within 12 months; verified protection in the approved target state for the function within 24 months, read as for Tier 1
Tier 3 – Standard	Any applicable risk dimension (1, 2 or 4) Medium or above	Begin within 24 months; complete within regulatory deadline
Tier 4 – Long Tail	Every applicable risk dimension Low. A blocked	Monitor; compensating controls; migrate when

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

Tier	Condition (the first that matches assigns the tier)	Target Timeline
	migration does not place an entry here: the entry keeps the tier the test gives it and carries the delivery state blocked	feasible

Sequencing within a tier. The composite score orders the entries inside each tier. Tier 1 in a large estate is hundreds of entries, and the score is the position each takes within it, read with the Track A and Track B sequencing logic of Activity 3.3: the track rules sequence the exposure classes within the tier, and the score ranks the entries inside each class. The score is an uncalibrated ordinal construct. It ranks entries against each other, it is not calibrated to a probability, a loss or a date, and no score band is published for it or read as one.

Dimension 3 reaches the program through the score and the roadmap, not through the tier: a Hard entry sequences earlier inside its tier because of its lead time, and its delivery state and firmness grade (Activity 7.5) set where it appears on the roadmap (Activity 4.2).

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:

- **Track A – key exchange (TLS, VPN, key establishment protocols).** Sequence by exposure, in three steps:
    1. Internet-facing TLS and partner VPNs. Each channel migrated here stops adding sessions to whatever a positioned collector is recording.
    2. Internal east-west traffic.
    3. Application-embedded key establishment.
- **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, because a forged update propagates to every downstream system that trusts it; then CA modernization and key-lifetime reduction, then end-entity and document signing. TNFL exploitation requires an operational CRQC, but

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

signature migration has the longest lead times in the program (PKI architecture decisions, toolchain support, HSM dependencies), which is why Track B starts in parallel rather than waiting its turn.

Certificate authorities are blocking counterparties on this track: an issuer that cannot issue ML-DSA or FN-DSA certificates, or composite certificates, on the dates in the roadmap gates every migration behind it. Each CA's post-quantum issuance roadmap is collected through the certificate-authority questions in Activity 7.2.

- **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. Stores with the longest confidentiality horizons must not fall off the roadmap, and an exposed store (a cloud archive, a backup export, media that has left the organization's control) enters the earliest tranche on its own exposure.

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, and 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

APPLYING THIS PHASE

- **Who runs it.** The QRPM owns the scoring model and brings it to the SteerCo for ratification. The Cryptographic Record workstream runs the scoring against the CBOM. The second line validates a sample after each run, under the model-risk treatment in GRC Implementation. The QRPM is accountable for the QRA, and the SteerCo for the appeals decisions.
- **Minimum viable version for a small estate.** Use the four dimensions with the default weights, ratified by the security forum that acts for the SteerCo. Score Priority A entries first. Leave unknown exposure unscored under the two rules of Activity 3.1, and list it by owner. Assign tiers by the ordered test of Activity 3.2, and use the

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

composite score only to order entries inside a tier. Produce the QRA as a ten-page document with the tier table, the agility-debt register and the compliance mapping. Do not tune the weights before the first run has been read.

- **The decision this phase produces.** Ratification of the scoring model and the appeals rule before the first run, and acceptance of the tier assignments and the sequencing logic after it, both taken by the SteerCo at gate G2 → G3. An appeal against a tier assignment is decided under the pre-approved rule and minuted with the evidence.
- **What it hands on.** Phase 3 hands four things to later work:
  - The tiered backlog and the Track A and Track B sequences, to Phase 4 for the roadmap and to Phase 5 for pilot selection.
  - The agility-debt register, to the Crypto-Agility workstream and the Metrics foundation.
  - The QRA and its compliance mapping, to GRC Implementation and the evidence dossier.
  - The unscored-entry list, to Phase 1 and Phase 2, as gap-closure work.

OUTPUTS

Output	Quality Criteria
Scored and tiered CBOM	Every CBOM entry is either scored and tiered under the ratified model or listed in the unscored-entries report; no entry has a provisional or defaulted score
Quantum Readiness Assessment (QRA)	Executive summary, heatmap, prioritized backlog, gap analysis, compliance mapping
Migration sequencing recommendation	Ordered backlog ready for Phase 4 roadmap planning
Agility-debt register	Every Tier-1 and Tier-2 entry with Hard control posture or system age, or Critical verifier updatability, listed with reason and owner; feeds the Crypto-Agility workstream and the roadmap
Unscored-entries report	Count and age of entries flagged "classification missing", each with the missing field, the classification owner, its review priority (internet- and partner-facing entries first) and a resolution date; none older

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

Output	Quality Criteria
	than one quarter

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 set the start windows, and composite scores order the entries inside each tier. Migration feasibility ratings and delivery states identify the entries that wait on vendor support, and show where each one appears on the roadmap.
- **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.
- **Sequencing a tier as if every entry could start tomorrow.** Scheduling a blocked entry (no vendor support, no standard) ahead of the high-risk entries that could be migrated today. The tier stays where the risk puts it. The delivery state and the bridging pattern (Activity 7.4) set when the work can start, and inside a tier the available entries start first.
- **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.

Common failure: a backlog that reports classical remediation as PQC progress, because closing a TLS 1.0 finding and migrating a key exchange were counted on the same line.

CHALLENGE QUESTIONS

- Which entries are unscored, who owns them, and how old is the oldest?
- Which weights changed since ratification, under which minuted decision, and what moved tier as a result?
- Which Tier-1 assignments were appealed, and how was each decided under the appeals rule?

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

- What share of Tier 1 carries agility debt, and which entries are Critical on verifier updatability?
- Which findings closed this quarter were classical remediation, and were any reported as PQC progress?

MATURITY INDICATORS

Level	Indicator
Level 0 – Unaware	No quantum risk assessment exists
Level 1 – Aware	Informal awareness of which systems are "probably vulnerable"; no structured scoring
Level 2 – Initiated	Formal risk scoring model applied to Tier-1 CBOM entries; prioritized migration backlog exists; QRA document produced
Level 3 – Progressing	QRA updated quarterly; all CBOM entries scored and tiered; migration sequencing drives Phase 5 execution; legal risk dimension assessed
Level 4 – Advanced	Continuous risk posture management; automated re-scoring when CBOM changes or regulatory deadlines shift; QRA integrated into enterprise risk register and audit cycle
Level 5 – Optimized	Continuous risk posture management; re-scoring automated on CBOM and regulatory change; the agility-debt register at zero unaccepted debt; appeals rare and minuted

© 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) states what needs to happen and in what order. A roadmap adds the dates, the owners, the cost, and the dependencies. 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:

- **Organization:** Executive sponsor and QRPM confirmed. SteerCo calendar and charter published. RACI matrix for all workstreams.
- **People:** 10-to-20 person training cohort enrolled in PQC fundamentals; "crypto champions" designated per platform team.
- **Plan:** Crypto-BOM v1 structure defined and Priority A population underway, with v1 at ≥70% Tier-1 coverage publishing in Q2. Two hybrid pilots selected, TLS and VPN. Vendor questionnaires sent to the top 10 strategic vendors.
- **Policy/Procurement:** Cryptographic policy updated, so that approved cipher suites include PQC hybrids and new certificates carry shorter key lifetimes. PQC and crypto-agility clause added to all new RFPs and contract renewals.
- **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; initial board briefing delivered if Phase 0 has not already delivered it (Activity 0.3)	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 preliminary pilot targets (TLS + VPN) from the Priority A scoping list (Activity 1.0), to be confirmed against the Phase 3 tiers before any production canary; 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	Complete Priority A	Execute pilots in lab/staging;	Collect vendor responses;

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

Quarter	Governance & People	Discovery & CBOM	Pilots & Technical	Vendor & Policy
	review, ratifying the scoring model, its weights, the ordered tier test and the appeals rule before the first scoring run (Activity 3.1); first quarterly board KPI report against the baseline (the Activity 0.3 cadence); present initial findings to executive sponsor; refine RACI based on Q1 experience	discovery; begin Priority B; publish Crypto-BOM v1 (≥70% Tier-1 coverage); begin Phase 3 risk scoring on the populated Priority A entries and confirm the pilot targets against the resulting tiers; surface and report classical vulnerabilities	measure performance baselines; document findings and compatibility issues	classify vendors by PQC criticality; publish updated cryptographic policy and procurement language
Q3	Second SteerCo review; quarterly board KPI report; present QRA findings; request Year 2 budget refinement if needed; expand training to second cohort	Complete Phase 3 risk scoring on Priority A systems under the ratified model; populate CBOM migration status fields; integrate CBOM into CI/CD (initial)	Promote validated pilots, confirmed as Tier 1 by the Phase 3 run, 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 with Year 1 achievements and the Year 2 plan; finalize multi-year	Complete QRA for Tier-1 systems; Crypto-BOM v2 (≥90% Tier-1, ≥60% Tier-2); establish	Complete canary validation; begin Wave 2 (internal production); document validated migration	Vendor PQC scorecard published; escalation plan for non-compliant strategic

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

Quarter	Governance & People	Discovery & CBOM	Pilots & Technical	Vendor & Policy
	roadmap	continuous discovery operating rhythm	patterns as repeatable playbooks	vendors; Year 2 vendor engagement calendar set

The Crypto-Agility workstream carries four Year-1 milestones alongside the columns above: Q1, the workstream owner named and the first rehearsal target selected from the pilot candidates; Q2, the approved-algorithms policy published in machine-consumable form and consumed by at least one pilot; Q3, the first algorithm-change rehearsal executed in staging on a pilot system, timed across assessment, decision, change and verification (Activity 5.2); Q4, the agility-debt register baselined from the Tier-1 QRA and its Year-2 burn-down target 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
Year 4	Long Tail & OT	Tier-3/4 systems addressed; OT gateway protections deployed; embedded device migration or containment; compensating controls where

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

Year	Focus	Key Milestones
		migration infeasible
Year 5	Hardening & Agility	Eliminate remaining compensating controls; achieve crypto-agility across the estate; transition from hybrid to PQC-only where counterparties and vendors support it; achieve target maturity level

The Coverage KPI reads this roadmap against all Tier-1 endpoints, internet-facing and internal (Metrics, KPIs & Reporting). The Year-2 milestone puts every Tier-1 internet-facing endpoint on hybrid key exchange, which is why the Year-2 target is below 100 percent. The internal Tier-1 endpoints follow in Year 3 with the east-west rollout, which raises the figure to 95 percent. The KPI lags the roadmap by design. A Year-2 reading below target with the internet-facing milestone met means the internal Tier-1 population is larger than the plan assumed, which is a denominator finding and is reported as one.

The roadmap must show the Crypto-Agility workstream's milestones on the same plan as the migration waves. Agility built afterward is a second program.

The milestones follow the four-year sequence in the Crypto-Agility foundation:

- **Year 1.** Policy and rehearsal foundation (Activity 4.1).
- **Year 2.** Provider-pattern abstraction in Tier-1 services, as each wave migrates them, and the first rehearsal on a non-production Tier-1 replica.
- **Year 3.** A production rehearsal on a Tier-1 system per track, inside the agility window, with the configuration-driven share reported against its target.
- **Year 4 and beyond.** A standing rehearsal cadence, and agility-debt burn-down to zero unaccepted debt before closure.

Each year's configuration-driven coverage target is set in the roadmap and read by the agility KPIs (Metrics, KPIs & Reporting). Blocked Tier-1 entries (delivery state blocked, Activity 3.2) appear on the roadmap against their unblocking date, the vendor date with its firmness grade, with the bridging pattern that covers them until then. They are never moved to the long tail because the path is closed.

4.3 Align to Infrastructure Refresh Cycles

Map migration tasks to existing planned expenditures:

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

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 (validated PQC key types in current firmware; models are listed in Tool and Vendor References)
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 ten workstreams defined in Phase 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

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

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. A preliminary selection from the Priority A scoping list (Activity 1.0) may precede the first scoring run so that lab work starts in Q1 (Activity 4.1). The targets are confirmed against the Phase 3 tiers before promotion to a production canary, and a target the run does not place in Tier 1 is replaced
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.

Upstream Activity	Must Complete Before	Why
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 why vendor engagement must start early, before the full CBOM is complete. AI-assisted tooling (Activity 5.7) compresses engineering effort in triage, enrichment and drafting. It does not compress the vendor critical path, certification cycles or change windows, and a roadmap that books the tooling's speed against those segments is wrong by the length of the longest of them.

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

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

- If leading indicators worsen (e.g., CRQC timeline estimates shortened) → Move to accelerated track with pre-drafted resource request
- 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 the bridging pattern, and escalate contractually

4.6 Define Milestone Gates

Gates are of three kinds:

- **A rollout gate** gives permission to start work on a population. G4 → G5 approves the pilot designs, and G5 → G6 approves the Tier-1 rollout.
- **A wave-exit gate** takes the evidence needed to leave a wave (Activity 5.4).
- **The closure gate** closes the program (Migration Verification & Program Closure).

Each gate applies to the population named in its record, not to the whole estate. The identifiers below are unchanged from earlier editions:

Gate	Criteria	Decision Authority
G0 → G1	Charter approved; budget committed; QRPM appointed	Executive Sponsor
G1 → G2	≥70% Tier-1 inventory complete; CBOM structure defined; accepted-gaps register signed	SteerCo
G2 → G3	CBOM populated; risk scoring complete under the ratified model; QRA delivered; agility-debt register produced	SteerCo
G3 → G4	Multi-year roadmap approved, with the agility milestones on it; Year 1 plan resourced	Executive Sponsor
G4 → G5	Pilot designs approved; success criteria defined; rollback plans documented; first rehearsal target selected	SteerCo
G5 → G6	Pilots validated; performance	SteerCo

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

Gate	Criteria	Decision Authority
	baselines established; agility criterion applied to every pilot, with any agility debt recorded with a closure date; Tier-1 rollout approved	
G6 → G7	Infrastructure upgrades scheduled; vendor commitments tracked and graded for firmness (a scheduling and grading gate, held by the QRPM; a capital decision inside it goes to the Executive Sponsor under Phase 6)	QRPM
G7 → Closure	Closure criteria met (Migration Verification & Program Closure): verification dispositions for Tier 1 and Tier 2, the closure proof, zero unaccepted agility debt as a governance condition with the debt total reported in full, the closure report's exposed-services, compensating-controls, residual-risk and remediation-date lines, the BAU handover register	The governance level that chartered the program

The accepted-gaps register at G1 → G2 generalizes a device the Accelerated Execution Profile (Activity 4.7) already uses ("documented Layer 3 and Layer 4 gaps") to the normal path: a gate that passes on coverage alone passes on the tools' view of the estate, and the register is what states what the tools did not reach.

The agility criterion at every migration gate. From G5 → G6 onward, and at every wave gate in Activity 5.4, a system counts as migrated only when all four of these conditions hold:

1. The negotiated outcome is observed in production traffic (Migration Verification & Program Closure).

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

2. Classical-only negotiation is refused where policy requires it, and the refusal is tested.
3. The algorithm can be changed by configuration, without a code change or a rebuild.
4. The change has been rehearsed in staging, with the rollback exercised (Activity 5.2).

The first two conditions decide the protection result. The last two decide the agility result. A system that meets the first two and fails the last two has been migrated once and will be migrated again by hand at the next algorithm change. The gate must not pass it as migrated. It passes it as migrated with agility debt, recorded in the Phase 3 register with a closure date. A system that fails either of the first two has not passed protection and is not migrated at all. It stays in its wave with its exposure recorded, and it does not count as migrated with agility debt.

The four conditions are the test for negotiated protocols: TLS, IKE, SSH and any function whose algorithm is offered and accepted in a handshake. Six other cryptographic functions are not negotiated in a handshake:

- certificate-based authentication
- firmware, update and secure boot
- long-term document and evidence preservation
- stored data and wrapped keys
- application and account authorization
- proof systems and commitments

For each of these, read the first two conditions from the function's row of the verification matrix in Migration Verification & Program Closure. That row states the positive evidence and the negative or boundary test to use in place of observed negotiation and tested refusal. The third and fourth conditions apply unchanged.

A signed firmware image has no handshake to observe: its first condition is a production verifier that accepts the post-quantum signature, and its second is a sampled verifier that rejects an image carrying only a classical signature.

Two results in every gate record. A gate record from G5 → G6 onward shows two results as separate fields: the protection result (which functions and boundaries are verified, which remain exposed) and the agility result (which change was rehearsed, through which mechanism, in what elapsed time, with which constraints). A system can pass on protection and still have agility debt recorded against it, with a dated closure entry. Neither result stands in for the other, and the Coverage and Agility KPIs read them separately (Metrics, KPIs & Reporting).

© 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. A fourth trigger is silence: a laboratory or government the SteerCo designated in advance ceasing to publish cryptanalytic results, or gating their disclosure, is a reassessment event and never a reassurance.

Second, the compressed sequence. Four things change in it:

- 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.
- Lower-tier systems move directly to compensating controls (segmentation, overlay encryption, key-lifetime reduction) instead of 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.

4.8 Re-Baseline Protocol

Activity 4.7 pre-drafts the case for going faster. This activity pre-drafts the case for slipping, so that the program has two pre-approved speeds and neither is improvised. A re-baseline moves committed dates through a governed procedure, with the evidence that the adaptation was governed filed where an auditor will look for it.

- **Triggers.** Any one of the following must open a re-baseline review within the next SteerCo cycle: actual effort on a wave exceeding the estimate by a factor the SteerCo set in advance (1.5 is a workable default); a Strategic Blocking vendor

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

slipping a committed date beyond the bridging capacity recorded in Activity 7.4; a pilot failure of a class the pilot plan defined as blocking (an interoperability failure with a Tier-1 counterparty, a security finding in the deployed implementation); or movement of an external deadline the roadmap was built against.

- **Authority map.** Moving a date requires the authority that owns it. A wave-internal milestone may be moved by the QRPM with notice to the SteerCo. A gate date or a wave completion date requires SteerCo approval. A date the board has seen, or any date that consumes the regulatory compliance buffer in the risk appetite statement, requires the Executive Sponsor and board notification under the escalation triggers in Activity 0.3.
- **Evidence rule.** The re-baseline is minuted with the trigger, the options considered, the decision and the new dates, and the minute is filed to the evidence dossier. A minuted re-baseline is proof that the program adapted under governance. An undocumented slip reads, to an auditor or a regulator, as a program that lost control of its plan.
- **Cadence.** Programs with a pre-built procedure re-baseline in days. Programs that improvise re-baseline in quarters, because the decision waits for the next committee that believes it has the authority to take it.

APPLYING THIS PHASE

- **Who runs it.** The QRPM owns the roadmap and the PMO. The PMO runs the gates and the SteerCo cadence. The Crypto-Agility workstream owner places the agility milestones on the plan. The Executive Sponsor approves the multi-year roadmap at gate G3 → G4, and holds the activation authority for the Accelerated Execution Profile.
- **Minimum viable version for a small estate.** The Year-1 plan on a two-quarter rhythm inside the merged charter-plus-plan; three gates (G1 → G2, G3 → G4, G5 → G6) with the agility criterion at the last; the refresh-cycle alignment table of Activity 4.3 built from the existing asset-refresh calendar; the Accelerated Execution Profile and the Re-Baseline Protocol each as a one-page pre-approved procedure. A PMO is a weekly meeting and a tracked plan, not a function.
- **The decision this phase produces.** Approval of the multi-year roadmap with the agility milestones on it and the Year-1 plan resourced, taken by the Executive Sponsor at gate G3 → G4. Every later date change is taken under the authority map of Activity 4.8 and minuted.
- **What it hands on.** Phase 4 hands five things to later work:
  - The wave plan, the gate criteria and the pilot targets, to Phase 5.
  - The infrastructure budget and the refresh alignment, to Phase 6.
  - The vendor dependency dates and the re-baseline triggers, to Phase 7.
  - The KPI baselines and targets, to the Metrics foundation and the board report.

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

- Every gate record and re-baseline minute, to the evidence dossier.

OUTPUTS

Output	Quality Criteria
Multi-year roadmap	5-year plan with annual milestones and the agility milestones; aligned to refresh cycles; critical path identified
Re-Baseline Protocol and Accelerated Execution Profile	Triggers, authority map and evidence rule approved in advance; both exercised annually as tabletop scenarios
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

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.

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

- **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.

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.

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

- **Governance without teeth.** Establishing a SteerCo that meets but does not make binding decisions, approve funding, or hold workstream leads accountable. The SteerCo must have decision authority over budget, timelines, risk acceptance, and vendor escalation.

Common failure: a roadmap that moves its dates without a minuted re-baseline, so that the plan an auditor reads and the plan the program runs are two different documents.

CHALLENGE QUESTIONS

- Which committed dates moved this period, and where is the minuted re-baseline for each?
- Which gates were passed with an open criterion, and which criterion was waived by whom?
- Are the agility milestones on the same plan as the migration waves, with owners and dates?
- Which SteerCo decisions this quarter changed a date, a scope or a risk acceptance?
- Has the Accelerated Execution Profile been exercised in the last twelve months, and who can invoke it today?

MATURITY INDICATORS

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

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

Level	Indicator
	execution to ongoing posture management
Level 5 – Optimized	The roadmap has become ongoing cryptographic posture management; both pre-approved speeds exercised annually; every date change is recorded in a minuted re-baseline

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

PHASE 5 – PILOTS & MIGRATION EXECUTION

PURPOSE

Execute the migration through controlled pilots, validate the patterns, and scale to production in 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, which become Phase 6 requirements. Infrastructure upgrades then enable broader pilot scope. And 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. By this point the skills gap assessment from the Skills & Team Structure foundation should have been addressed. 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. Where a mirrored environment cannot be built, the small-estate profile in Applying This Phase states the one documented exception. A production change window is never treated as a test environment.
- **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

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

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.

DEPLOYMENT ENVIRONMENT CLASSIFICATION

Understanding Deployment Constraints

Before selecting pilot targets or designing hybrid deployments, organizations must classify their systems by deployment environment, because that 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 September 2026, the CMVP validated-modules list carries post-quantum operations as approved services in software and in hardware, and the two certificates this framework cites are not interchangeable evidence.

Property	Certificate #5247	Certificate #5497
Module	Go Cryptographic Module (versions v1.0.0 and v1.0.1, Geomys LLC)	QASM Cryptographic Module (Crypto4A Technologies, hardware version QASM 1.1, firmware 5.0.1.158)
Type and level	FIPS 140-3 Level 1 software module	FIPS 140-3 Level 3 multi-chip standalone hardware module
Dates	Validated April 27, 2026; updated July 7, 2026	Initially validated August 19, 2026; updated August 21, 2026; sunset date August 18, 2031
ML-KEM	KeyGen and EncapDecap, which covers encapsulation and decapsulation, at ML-KEM-768 and ML-KEM-1024 only, not ML-KEM-512	KeyGen and EncapDecap at ML-KEM-512, -768 and -1024
ML-DSA	None. The Go module version that implements ML-DSA	KeyGen, SigGen and SigVer at ML-DSA-44, -65 and -87

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

Property	Certificate #5247	Certificate #5497
	(v1.26.0) is on the Modules in Process list, not on the validated list	
Other approved algorithms	Not applicable	SLH-DSA across all twelve parameter sets, and LMS

Every citation of these certificates states the qualifications below.

- **#5247.** ML-KEM EncapDecap reaches the approved Key Transport service through the Key Encapsulation security function, and the Security Policy contains no statement declining KEM-based key establishment. Two qualifications apply to this citation. The approved-services table names services generically, and reaches ML-KEM only through the security-function table. And §2.10 states that the module supplies cryptographic services at the API level to a calling application, and does not implement full key agreement inside its boundary.
- **#5497.** Two qualifications from its Security Policy apply to every citation. The first is that §2.7 states that the module does not establish SSPs using an approved KEM and instead exposes ML-KEM functionality for an external application to use as part of an approved KEM. The second is that HSS key generation, signing and verification are vendor-affirmed, not CAVP-tested, with HSS used alongside ECDSA in the module's firmware-update verification service.

A sentence asserting that approved post-quantum key establishment exists in a validated module cites one certificate with its qualification (#5247 with its parameter-set limit, #5497 with its §2.7 statement), never both as a single proposition. Every ML-DSA statement in this framework rests on #5497 alone. Some validated modules still list PQC only among their non-approved services (AWS-LC among them), and a listing of that kind does not count.

The planning rule is per product, not per calendar. For each deployment, record:

- the governing requirement
- the exact module and version
- the certificate number and status
- the approved services the certificate lists
- the required security level
- the tested operating environment

Confirm that the certificate covers the cryptographic operation the deployment needs. An algorithm standard, a CAVP algorithm certificate, a product's "FIPS mode" and a module's presence in the queue are each distinct from a completed validation covering

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

the operation. Where no suitable validated product exists for a system, the dependency is recorded with the vendor's evidence, the interim controls, the exception authority and the next decision date. It constrains that migration and does not set a date for the program.

Two Implementation Guidance updates in 2026 bear on planning:

- **April 9, 2026.** Added the SP 800-227 testing scenario for hybrid KEMs (IG D.S, Scenario 2), so a module can be validated with PQ/T hybrid key establishment in its approved boundary. It also added the entropy requirement stated in Activity 6.2: physical noise sources evaluated at multiple operating conditions, strongly recommended through December 31, 2026 and required afterward (IG D.K, Resolution 23).
- **August 19, 2026.** Added IG 1.B on hybrid modules, and the word means something different there: a hybrid module in FIPS 140-3 is a software or firmware module with a hardware component, not a PQ/T hybrid algorithm.

A vendor statement that a module is "hybrid-validated" must be read against that definition. Three points for practitioners. FIPS mode is a configuration. Validation is a certificate, tied to a module version and a tested operating environment. The deployment environment classification keys on the certificate, never on the configuration.

FIPS 140-2 certificates moved to the Historical list on September 21, 2026, the date stated in the maintained prose of NIST's FIPS 140-3 Transition Effort page (its transition-schedule table, stamped 2021, prints September 22, 2026). CMVP states a permission for those modules, not a prohibition:

- It supports their purchase and use for existing systems and notes that the selection of FIPS 140-3 modules may be limited for several years.
- It recommends that purchasers consider every module on the validated-modules search page regardless of whether it was validated against FIPS 140-2 or FIPS 140-3.

The exclusion of a Historical module from procurement for new systems follows from that scoping to existing systems and from agency and contractual policy, not from a CMVP prohibition. Whether an existing deployment on a historical certificate may continue is decided under the applicable agency or contractual risk policy, and the two questions are assessed separately.

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) ship the NIST-standardized algorithms and hybrid ML-KEM key exchange in TLS 1.3. PQC Deployment Status, below, states what "by

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

default" covers for each. 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, protection profile or policy that mandates FIPS 140-3 validated cryptographic modules: federal government systems, defense contractor environments handling CUI under an obligation that names validated modules, and any system where a contract or an organizational policy requires them. The requirement is checked at its source, not assumed from the sector. HIPAA's encryption specifications are addressable and do not by themselves mandate FIPS validation. A payment HSM's obligation comes from PCI PTS as much as from FIPS 140-3.

PQC enters production in these environments only through a module whose certificate lists the operation as an approved service. Validated software and validated Level 3 hardware both exist (above), each covering the operations and parameter sets its certificate names, and any other product is read from its own certificate. Plan pilot and staging work now. Target production deployment per system on the certificate date of the module that covers its operation, with the dependency recorded until then. 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. Under the 2024 update to CNSSP 15, new NSS acquisitions must be CNSA 2.0 compliant from January 1, 2027; equipment and services that cannot support CNSA 2.0 are phased out by December 31, 2030; and CNSA 2.0 algorithms are mandated for use by December 31, 2031, each unless otherwise noted in the applicable profile.

The advisory's per-category chart sets its own dates:

Category	Chart date
Software and firmware signing	2030
Networking equipment	2030

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

Category	Chart date
Web, cloud and operating-system platforms	2033
Niche equipment	2033
Updating or replacing custom applications and legacy equipment	2033

The chart gives no date beyond 2033, and its 2033 dates fall after the CNSSP 15 mandated-use date of December 31, 2031. The chart does not extend the policy: a system's binding date is the CNSSP 15 date or the date its applicable profile notes, and the chart date is recorded with its provenance, as the advisory's planning guidance for that category.

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.

A FIPS 140-3 validation that covers post-quantum algorithms, at Level 3 or any level, is not by itself acceptability for a national security system: a certificate covers an operation on a module version in a tested environment and nothing beyond it, the discipline this classification applies to every environment class, and CNSA 2.0 acceptability is read separately, operation by operation, against the algorithm list in Appendix A.

Certificate #5497, the Level 3 hardware certificate this framework cites, shows why. Its approved algorithms include:

- **ML-KEM and ML-DSA**, which CNSA 2.0 requires at ML-KEM-1024 and ML-DSA-87.
- **SLH-DSA at every parameter set**, which CNSA 2.0 allows for no NSS use.
- **LMS**, which CNSA 2.0 allows only for signing firmware and software (Appendix A).

The module verifies its own firmware updates with ECDSA and with HSS, and HSS key generation, signing and verification are vendor-affirmed rather than CAVP-tested. CNSA 2.0 does not allow HSS for NSS. A reader who takes the Level 3 validation as NSS acceptance is wrong for two of the module's own algorithms.

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, the date is the certificate date

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

of the module that covers the operation, recorded per system. 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.

TWO-TRACK MIGRATION MODEL

Track A (Confidentiality) and Track B (Integrity and Authentication)

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. National mandates now set separate deadlines for key establishment and for signatures, which institutionalizes the split this framework has treated as foundational since v2.0: Executive Order 14412's high-value-asset dates are December 31, 2030 for key establishment and December 31, 2031 for digital signatures, and the published migration roadmaps of the major cloud providers sequence the two tracks the same way.

OMB Memorandum M-26-15 (June 24, 2026) executes the order for federal agencies outside national security systems: each agency submits its PQC Migration Plan to OMB and the Office of the National Cyber Director by October 22, 2026, and the memorandum's five phases run prioritized key-establishment migration through 2030, signature migration in 2031 and full migration in 2035.

Track A: Confidentiality (Key Exchange) addresses Harvest Now, Decrypt Later. This framework treats HNDL as an active planning risk and states the basis for that treatment, because the harvest itself cannot be observed: passive interception leaves no trace in the victim's network, and no recorded session announces that it was recorded. The treatment has four grounds. Means: passive optical and radio interception at scale is demonstrated capability, and the storage it requires is now cheap.

Motive: the target data keeps its value for decades, longer than any credible CRQC estimate. Opportunity: backbone and landing-station positions, and physical proximity to facilities, are held by the actors with the means. Precedent: seven decades of documented cable and backbone interception, from Cold War landline taps to the collection programs disclosed since 2013. A rational actor with those four does not wait for the decryption capability before collecting, and the framework plans on that actor.

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

An adversary who captures ciphertext today keeps it, so every day that key exchange remains quantum-vulnerable creates permanent exposure for data with long confidentiality requirements. The mitigation is hybrid key exchange – X25519MLKEM768 on TLS channels, and the RFC 9370 multiple-key-exchange profile with ML-KEM on IPsec and VPN channels, wherever they carry sensitive data – and this track has the highest urgency for that reason.

Corrected at v3.0: earlier editions stated harvesting as an observed fact in this paragraph, in SOC Use Case 5 and in Appendix F. The position is inference from demonstrated capability. The planning posture is unchanged. The description of its basis is corrected, and the Version History records the change. Appendix F handles the objection that no harvesting has been shown.

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 depends on quantum-vulnerable signatures that a certificate-centric view misses. FIDO2/WebAuthn passkeys typically assert with ES256 or RS256 (P-256 and RSA, both Shor-vulnerable). Inventory the algorithms actually enrolled. IANA registered PQC COSE algorithm identifiers in 2025 and ML-DSA for JOSE and COSE is RFC 9964 (May 2026); IETF work on ML-DSA in WebAuthn is active, but deployable post-quantum passkeys depend on authenticator and platform support that is still uneven.

OAuth and OIDC tokens are signed with RS256 or ES256, and SAML assertions with the RSA and ECDSA algorithms of XML Signature. 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 change.

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

Treat short-lived session tokens as low priority, because their TNFL exploitation window is bounded by their lifetime. The token-signing key that a CRQC would recover mints new credentials with fresh expiries, so the Track B priority attaches to the key and not to the individual token. 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 and conservative. CNSA 2.0 recommends them for signing software and firmware. This is the signature migration NSA prefers to begin now, because validated implementations of those schemes are commercially available.

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

ML-DSA-87 is approved under CNSA 2.0 for all signature use cases, including firmware and software signing. NSA names it as possibly the more appropriate choice in two situations: where a signing strategy needs more signatures than one LMS or XMSS key can reasonably produce, and in distributed signing environments (NSA CNSA 2.0 FAQ, Version 2.1, December 2024).

What the validated population supports:

- The validated LMS implementation this framework can evidence is CMVP certificate #5497, whose CAVP entry covers one parameter set, LMS_SHA256_M24_H5 with LMOTS_SHA256_N24_W2, while the module's own security-function table lists more.
- XMSS support is materially rarer than LMS, and the CMVP search form offers no filter for either scheme, so the validated population cannot be enumerated and no broader claim is made in either direction.
- The deploying organization verifies parameter-set coverage on the specific certificate rather than inferring it from the algorithm's presence.

"Deploy now" comes with SP 800-208's own conditions:

- Key generation and signature generation happen inside a hardware module with at least Level 3 physical security.
- The private key is never exported, in any form.
- Signing state is managed inside the module, so that backup, replication and failover designs cannot reuse a state.

A deployment that does not meet section 8 of SP 800-208 is not a conforming deployment, and the verifiers must accept the parameter set chosen.

State safety is an acceptance item evidenced inside the conforming implementation, not an inference from "HSM-backed" or "controlled ceremony":

- the supported parameter sets
- exhaustion monitoring
- atomic state handling
- tested backup, replication, crash-recovery and concurrent-signing behavior

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 ML-DSA toolchains and validated modules: 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.

HSM-backed signing plans on ML-DSA and SP 800-208. FN-DSA (FIPS 206) is under development at this edition, with no published draft and no published date (Appendix

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

A). Its signing depends on floating-point arithmetic that must be implemented in constant time, which creates side-channel surface that validated-module support has been slow to cover, so FN-DSA is the constrained-environment option after standardization and validation, not the HSM default. A Track B plan that assumes FN-DSA in the HSM before a validated firmware exists sets a date it cannot keep.

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 Status (data as of 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 (hybrid ML-KEM key exchange available in TLS 1.3):** OpenSSL 3.5.0+, Go 1.24+, Node.js 24.5.0+, BoringSSL (current), rustls 0.23.22+. Three different things are often collapsed into "PQC enabled by default": the primitives being present in the library, the protocol integration (the hybrid group in the TLS 1.3 handshake), and whether that group is offered without configuration. The list states the second. Whether the hybrid group is offered without configuration or has to be switched on is a per-library, per-version setting, verified on the version in use, and Migration Verification & Program Closure counts the negotiated outcome, not the library's capability.
- **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, much of it through CDN and package defaults; capability that arrived that way is incidental and is not Coverage (Metrics, KPIs & Reporting). 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

5.2 Design Hybrid Deployments

First-layer selection by architecture pattern. Activity 3.3 sequences Track A by exposure. This decision picks the single layer that, updated once, protects the most traffic for the architecture in front of the team. The table gives the default first layer per pattern. The four steps below it handle estates that mix patterns.

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

Architecture pattern	First layer to migrate	Why there	What it does not cover
Corporate web through a SASE service or forward proxy	The proxy or SASE edge: hybrid key exchange on the client-facing and the origin-facing legs	One configuration change covers every user and every application behind the edge	Direct-to-origin traffic that bypasses the edge; application-layer encryption inside the tunnel
Cloud-native microservices with a service mesh	The mesh control plane: the mTLS group and cipher policy set once and rolled out by the sidecars	East-west traffic is the largest volume in most estates, and the mesh already owns its TLS configuration	Workloads outside the mesh; managed-service endpoints the provider terminates (Activity 7.7)
Remote-access VPN	The concentrator and its client software: hybrid IKEv2 (RFC 9370) or the product's hybrid TLS-VPN group	Every remote session passes one enforcement point, and client rollout follows the concentrator's policy	Split-tunnel traffic that never enters the tunnel
OT and SCADA	The OT/IT boundary gateway: PQC termination at the gateway, with the zone behind it unchanged (OT & Critical Infrastructure Extension)	Devices inside the zone cannot be changed on the program's timeline. The gateway can	Any path that crosses the zone boundary without the gateway; firmware-signing integrity (Track B)
Email	Transport first (SMTP over TLS on the mail gateways), end-to-end second (S/MIME and OpenPGP, Appendix H)	Transport covers every message with one change on the gateways; end-to-end depends on every client and every recipient key	Messages relayed through a party that does not offer hybrid TLS; stored mail (Activity 5.6)
API-driven financial services	The API gateway and its mutual TLS: hybrid key exchange on the gateway,	Counterparties reach one enforcement point, and the gateway holds the	Direct connections that bypass the gateway; the signatures on the

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

Architecture pattern	First layer to migrate	Why there	What it does not cover
	client certificates under the private PKI (Phase 6)	negotiation policy	tokens and messages behind it (Track B)

Decide this per data flow, in four steps:

1. Map the protocol stack the flow crosses, from the client to the last hop that decrypts it.
2. Identify the most centralized point at which one change covers the most sessions.
3. Check vendor readiness at that point. If the product cannot negotiate hybrid key exchange at its current release, move to the next-best point instead of waiting for a roadmap date.
4. Deploy hybrid there, record the negotiated outcome, and move to the next flow.

The authentication tax. One layer handles confidentiality for a flow, and its migration is what the four steps decide. Every public-key authentication point on the same flow (the server certificate at the edge, the client certificates in the mesh, the token signatures behind the gateway, the firmware signatures inside the OT zone) stays an independent Track B surface with its own longer lead time, and none of them is covered by the confidentiality change. The CBOM entry for the flow records both.

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:

- Provides immediate quantum resistance without waiting for universal PQC adoption
- 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
- Reduces deployment risk – protected against both known quantum threats AND unknown weaknesses in new PQC algorithms

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

- 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

The registry rule. The hybrid groups are on the standards track: RFC 10024 (August 2026, Proposed Standard) registers X25519MLKEM768 with Recommended=Y in the IANA TLS Supported Groups registry, and SecP256r1MLKEM768 and SecP384r1MLKEM1024, the NIST-curve combinations, with Recommended=N (in all three groups the FIPS approval comes from the ML-KEM component running in a validated module. The classical component is non-approved in X25519MLKEM768 and additionally approved in the NIST-curve groups, which matters only where classical-only operation must also remain approved).

The pure ML-KEM groups (MLKEM512, MLKEM768, MLKEM1024) are specified in draft-ietf-tls-mlkem, an Informational document in the RFC Editor queue at this edition, with every code point marked Recommended=N, which under RFC 9847 means the IETF has made no statement on suitability. Auditors and architects read the registry, not the RFC label, because the Recommended column is where the IETF recorded its preference. The guidance is unchanged: hybrid X25519MLKEM768 is the default wherever confidentiality must outlive the connection; pure ML-KEM is a compliance profile for regimes that prescribe it, adopted once the RFC has published (Appendix D).

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 bugs, and ML-DSA implementations are new.

Published analysis of early ML-DSA implementation flaws (Bernstein, June 2026) models, on the author's stated assumptions rather than on observed failure rates, that dropping the classical layer would increase the share of breakable signature keys by roughly an order of magnitude over 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.

The hedge exists only where the verifier enforces it. A composite signature is verified as one object. A dual-signed artifact hedges only if the required verifiers accept it when both signatures validate, and reject it when one signature has been stripped. A verifier that accepts either signature alone provides compatibility and no hedge. The verifier acceptance policy is recorded with the format. Where composite formats are not yet

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

supported, dual-sign in parallel under that policy. 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.

Hybrid deployment architecture per protocol:

Protocol	Classical Component	PQC Component	Standard/Draft	Production Ready?
TLS 1.3 Key Exchange	X25519	ML-KEM-768	RFC 9954 (hybrid key exchange framework, July 2026, Informational); RFC 10024 (X25519MLKEM768, SecP256r1MLKEM768, SecP384r1MLKEM1024; August 2026, Proposed Standard); pure ML-KEM groups in draft-ietf-tls-mlkem (Informational, RFC Editor queue at this edition). The IETF UTA draft on PQC for TLS-based applications is the reference for application teams	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	RFC 10042 (August 2026, Informational); OpenSSH 9.9+ (mlkem768x25519-sha256; default from OpenSSH 10.0)	Yes – available in current OpenSSH
S/MIME / Email	ECDH P-256	ML-KEM-768	Published standards: RFC 9881 (ML-DSA in X.509 certificates and CRLs) and RFC 9882 (ML-DSA in CMS), both October 2025; RFC	Standards published; early vendor support; not yet broadly deployed

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

Protocol	Classical Component	PQC Component	Standard/Draft	Production Ready?
			9814 (SLH-DSA in CMS); ML-KEM in CMS per the LAMPS profiles	
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, 2025); IETF drafts	Not yet – standards groundwork only; require vendor algorithm agility
Token signing (JWT / OIDC / SAML)	RS256 / ES256	ML-DSA	RFC 9964 (ML-DSA for JOSE and COSE, May 2026, Proposed Standard)	Early – inventory signing keys and rotation now

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:** Open-source library with a FIPS 140-3 validated module. Read the certificate's approved-services list for the PQC operation required before relying on it in a validated environment. A module can list PQC only as non-approved services (Deployment Environment Classification), and validation is per module version and tested environment.

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

- **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.

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 hybrid PQ/traditional schemes; ANSSI makes hybridation mandatory for products seeking its security visa (Appendix D)
- **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
- **OMB M-26-15 (US federal agencies, June 24, 2026):** Describes hybrid architecture as a useful tool for managing risk during migration and as an "intricate and resource-intensive stopgap" that introduces its own risks and complexities; agencies evaluate its tradeoffs thoroughly and may implement it on their own risk assessments and technical requirements
- **NSA CNSA 2.0:** Accepts hybrid as interim only

Recommendation for multinationals: Adopt hybrid as the default deployment pattern (satisfying the strictest positions, the continental European recommendations and the ANSSI security-visa requirement) with the understanding that hybrid is an interim bridge to PQC-only as vendor and browser support matures. This "highest common denominator" strategy satisfies the jurisdictions that recommend or prefer hybrid and is accepted where hybrid is neither required nor prohibited, including under OMB M-26-15, which leaves the choice to the agency's own risk assessment. It is a default confirmed per deployment profile against the applicable authority (Appendix D), not a compliance conclusion for every jurisdiction. This framework keeps the default despite OMB's more skeptical position, on the defense-in-depth grounds stated above.

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

Long-lived sessions, resumption and KeyUpdate. A connection established with classical key exchange and kept alive is not protected by a hybrid capability enabled later. Resumption is protected by mode, not by the age of the ticket: a resumed handshake that includes a fresh hybrid key share (psk_dhe_ke with a hybrid group) protects the new application traffic even when the PSK came from a classical handshake, because both key shares feed the derived secrets; PSK-only resumption (psk_ke) and 0-RTT early data derive from the classical PSK alone and are not protected.

TLS 1.3 KeyUpdate derives new traffic secrets from the existing ones: it stays inside the traffic-secret lineage of the original handshake and creates no independent post-quantum secret. PSKs, resumption and KeyUpdate are therefore not interchangeable forms of rekeying (ePrint 2026/1703, August 2026). Pilot design for long-lived channels (database replication links, message-queue connections, VPN tunnels, service-mesh sidecars) must force a fresh handshake on a defined interval after hybrid enablement, must require the hybrid group on resumption and disable 0-RTT for the flows in scope, and may expire pre-migration tickets as the simpler policy.

Migration Verification & Program Closure counts handshakes in which the hybrid group was negotiated, fresh or resumed, not connections that happen to be open on a hybrid-capable endpoint.

- **Pre-standard artifact retirement.** Programs that piloted in 2023 and 2024 on draft identifiers (the X25519Kyber768Draft00 and SecP256r1Kyber768Draft00 code points, and the draft hybrid names used in libraries, load balancers and proxies of that period) must retire them on a dated schedule: RFC 10024 obsoletes the pre-standard Kyber768 code points and marks them Recommended=D in the IANA registry. Retirement is verified by a negative test that no draft group is offered or negotiated on any production endpoint, and the test is repeated in Migration Verification & Program Closure.
- **The rehearsal pattern.** Every wave must include an algorithm-change rehearsal in staging on at least one system in the wave: a parameter-set change (ML-KEM-768 to ML-KEM-1024 on a key-establishment path), a hybrid-to-composite switch on a signing path, or a certificate-signature-algorithm change on an internal issuer. The rehearsal is timed across four segments: assessment, from the trigger to the decision that a change is needed, using the 48-hour assessment test in the Crypto-Agility foundation; decision, approval under the standing change authority; change, the configuration or policy rollout to the target; and verification, observed negotiation of the new algorithm in staging traffic, with the rollback exercised.

The four elapsed times are recorded and feed the rehearsed time-to-change KPI (Metrics, KPIs & Reporting), and the segment that dominates the total is where to invest next. A wave whose rehearsal fails, or whose rollback fails, has found agility debt (Phase 3 register) before production did. For custody and HSM-backed signing, the

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

algorithm-change rehearsal is one of the three named above. A quarterly failover to the backup HSM partition exercises state management, failover and the SP 800-208 state design, changes no algorithm, and is recorded and reported as a failover exercise whose timings do not feed the rehearsed time-to-change KPI.

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 about 2.3 KB of additional key-exchange data per handshake: the client key share grows from 32 to 1,216 bytes and the server key share from 32 to 1,120 bytes (RFC 10024), adding 2,272 bytes across the two directions (2,336 bytes of hybrid key shares in total) before certificates and transport overhead. 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 raw for ECDSA P-256 (a few bytes more in the DER encoding carried in certificates and TLS). Signature size 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

5.4 Scale from Pilot to Production Through Waves

Wave	Scope	Prerequisite
Wave 0 – Lab/Staging	Isolated test environments	Pilot design approved

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

Wave	Scope	Prerequisite
	only	
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

Every wave exit from Wave 1 onward applies the agility criterion of Activity 4.6 to the systems in the wave: negotiated outcome observed, classical-only refused and tested where policy requires it, algorithm changeable by configuration, and the change rehearsed with rollback. For a system whose function is not a negotiated protocol (a signed firmware image in Wave 5, a document-signing service, a wrapped key store) the first two conditions are read from the function's row of the verification matrix (Migration Verification & Program Closure).

The wave's rehearsal record (Activity 5.2) is the evidence for the last two conditions. A system that fails them exits the wave as migrated with agility debt, not as migrated. The wave-exit record shows the protection result and the agility result as separate fields for every system in the wave.

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.

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

- **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.
- **AES-256 as the default for new and long-horizon deployments:** Use AES-256 for new deployments and for data with long confidentiality horizons, and where a profile requires it. AES-128 remains approved, and symmetric algorithms do not face the transition problem that Shor-vulnerable public-key cryptography does. A larger key does not repair a weak mode, nonce discipline, entropy or key management, which stay in scope on their own terms.
- **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. Lifetime reduction bounds the use of a compromised leaf key and builds rotation discipline. It does not make the signature algorithm quantum-resistant, and it does nothing for the issuing key that a CRQC would recover to mint new certificates. That key is Track B's work.

Confidential Computing as complementary protection: For environments processing highly sensitive data, Confidential Computing (hardware-based Trusted Execution Environments: Intel SGX and TDX, AMD SEV-SNP, Arm CCA) protects data in use, a gap that PQC alone does not address. The boundary is the enclave or the confidential virtual machine, and the adversary it defends against is one who controls the host software stack outside that boundary, the hypervisor and the platform operator included.

It does not protect data before it enters the boundary or after it leaves, and a flaw in the workload running inside it is outside the guarantee. Its guarantee rests on dependencies the program records and migrates as its own: remote attestation, verified before any secret is provisioned, whose evidence is signed with the vendor's attestation keys; the key-provisioning path into the enclave, which is a key-establishment flow (Track A); and the vendor's attestation service, firmware and microcode, which are a vendor dependency under Phase 7.

Where those hold, PQC protects data in transit, AES-256 protects data at rest and Confidential Computing protects data during processing, and combining them is defense in depth against both quantum and classical platform-compromise threats within the boundary each one draws.

- **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), re-wrap the existing data-encryption keys under a key-encryption key derived from a PQC KEM shared secret (a KEM establishes a

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

secret; the wrapping itself is an authenticated symmetric operation, per SP 800-227).

This protects the keys held under the organization's control without requiring data re-encryption, reducing the scope of the data-at-rest migration challenge. It does not protect a copy of the ciphertext and the old key envelope that an adversary has already taken, which Activity 5.6 records as historical exposure.

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 still in custody, not a copy of the ciphertext and its key envelope an adversary already holds (recorded as historical exposure, below); requires KMS/HSM PQC support; pairs with the key-wrapping pattern in Activity 5.5
Crypto-shred (destroy the keys)	Data with no remaining business use whose confidentiality need still runs, held on media or under a rule that cannot yet delete it outright (immutable backups, a retention period not yet expired), where counsel confirms that the retention obligation and any preservation period do not	Only as strong as the certainty that no key copies exist; produce destruction evidence per the closure standard

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

Strategy	When it fits	What to watch
	require the data to remain readable	
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 (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.

Neither re-encryption nor re-wrapping repairs a copy an adversary already holds. A copy of the ciphertext with its key envelope that left the organization's control before the change (an exfiltrated backup, a lost tape, a replica at a former provider) stays decryptable by a future CRQC whatever is done to the copy still in custody. It is recorded against the store as historical exposure, with the date range, the data category and its remaining confidentiality horizon. It is recorded in the verification record's residual-exposure column (Migration Verification & Program Closure), and it enters the closure report as accepted residual risk, or as an open obligation where an external duty attaches to the data.

PGP-encrypted archives are a long-horizon at-rest class of their own. An archive encrypted to an RSA or ECC OpenPGP key stays decryptable by a future CRQC for as long as it is retained, and re-encryption under the composite algorithms of RFC 9980 (Appendix H) requires re-keying every recipient, because the RFC defines no RSA composite. Inventory them on the backups line, with the recipient keys, and decide per archive among the five strategies above.

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

5.7 AI-Assisted Migration Under the Existing Review Gate

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 some of the remaining failures reach production, 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 were produced with AI assistance. The failure mode to refuse: treating AI throughput as a reason to compress review. The tooling saves time in the steps around review and none inside it.

Effort compression is not schedule compression: the tooling shortens the engineering steps, while the organizational effort of building the CBOM, the governance and the agility capability does not compress with them, and a plan that reads a tooling gain as a program gain re-baselines within two quarters (Activity 4.8). A migration whose automation depends on a single proprietary hosted model stack has taken a dependency that can be withdrawn, repriced or restricted by jurisdiction, and the plan prices that revocation risk the way it prices any other vendor dependency (Activity 7.7; algorithm sovereignty in the Regulatory & Standards Alignment Map).

APPLYING THIS PHASE

- **Who runs it.** The Network & TLS/VPN, PKI & Code Signing and Applications & Platforms workstreams execute the pilots and waves, with one lead per track. The Crypto-Agility workstream owner co-signs every rehearsal record. The QRPM is accountable for the wave, and the application owner for each system in it.

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

- **Minimum viable version for a small estate.** Two pilots, one per track: hybrid X25519MLKEM768 on the internet-facing edge, and certificate-lifetime reduction with a test dual-stack issuer on the internal PKI. Waves 1 and 2 combine and the rehearsal runs on one system. Where the estate cannot fund a test environment that mirrors production, the pilot is run as a documented limited-production procedure, which is a recorded exception to the Phase 5 prerequisite and not a cheaper form of it:
  - a blast radius stated in advance (one endpoint or one canary share of traffic, with the flows and counterparties it can affect listed)
  - an observation period in which the negotiated outcome and the error rate are watched from the first session
  - a rollback rehearsed before the change and executable by the named incident commander without further approval
  - the exception approved by name by the risk owner and minuted with the pilot design at gate G4 → G5

A production maintenance window is a change window. It is not a test environment, and no record describes it as one.

- **The decision this phase produces.** The go or no-go for each wave against the wave exit criteria of Activity 5.4, taken by the SteerCo at the gate (Activity 4.6); rollback during a wave is taken by the incident commander named in the rollback procedure without further approval.
- **What it hands on.** Phase 5 hands three things to later work:
  - Pilot results and the wave plan, to Phase 6 as infrastructure findings and to Phase 7 as vendor gaps, with the evidence attached.
  - The updated CBOM entry per migrated system, to the Cryptographic Record.
  - The rehearsal record with its four elapsed times, to the evidence dossier and the agility KPIs.

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

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

Output	Quality Criteria
	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

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 the failing middleboxes, the HSMs lacking PQC key capability, the network paths that cannot handle hybrid handshake sizes, and the PKI components due for 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.

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

Runs in parallel with

- **Phase 6** – Infrastructure modernization and pilot execution are iterative. Each pilot exposes an infrastructure gap. Each gap closed widens the scope of the pilots that follow. 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.
- **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.

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

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

UK, CNSA 2.0) require pure PQC. Plan the hybrid-to-PQC-only transition from the beginning.

Common failure: capability enabled by a package default, reported as migration before the negotiated outcome was observed and before classical-only was refused where policy requires it.

CHALLENGE QUESTIONS

The SteerCo asks these before accepting the phase output. The assurance questions in GRC Implementation apply the same tests.

- Which migrated systems show observed hybrid negotiation in production traffic, and which are marked migrated on configuration alone?
- Which migrated systems have rehearsed an algorithm change, with the four elapsed times recorded, and which carry agility debt?
- Which long-lived channels and resumption tickets in the wave were forced through a fresh hybrid handshake, and which are still running on a classical lineage?
- Which pilot targets came from the Tier-1 list, and which came from convenience?
- Where is the rollback record for each wave, and who executed it?
- Which endpoints reported as migrated carry a CDN or package default rather than an enforced policy?

MATURITY INDICATORS

Level	Indicator
Level 0 – Unaware	No PQC pilots planned or underway
Level 1 – Aware	Lab testing only; no production exposure
Level 2 – Initiated	2+ production pilots running with measured results; rollback procedures tested; validated patterns documented
Level 3 – Progressing	Tier-1 internet-facing systems on hybrid/PQC, verified in use under the enforcement counting rule; wave rollout underway for Tier-2; defense-in-depth measures (tokenization, AES-256, segmentation) deployed
Level 4 – Advanced	Estate-wide hybrid/PQC deployment

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

Level	Indicator
	substantially complete; transitioning selected systems from hybrid to PQC-only; algorithm change rehearsed on Tier-1 systems inside the agility window
Level 5 – Optimized	Algorithm changes are routine operations executed inside the agility window on a standing rehearsal cadence; PQC-only wherever counterparties and vendors permit

© 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. A hybrid handshake sends the larger key shares. A post-quantum certificate chain sends the larger signatures. Together they raise 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 runs iteratively alongside Phase 5 and not as a discrete phase after the pilots: each early pilot reveals which infrastructure components need modernization, and each upgrade makes broader and more ambitious ones possible. 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

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

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 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.

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 which devices are 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.

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

- **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.

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. RFC 9881 (October 2025) is the published certificate and CRL profile for ML-DSA, and RFC 9882 the CMS profile. The trade-off is size: ML-DSA-65 signatures are approximately 3,309 bytes versus 64 bytes raw for ECDSA P-256 (a few bytes more in the DER encoding carried in certificates and in TLS).

In a scenario chain of five ML-DSA-65 signatures and two ML-DSA-65 public keys, built for illustration and not a typical chain, the authentication overhead in a TLS handshake would exceed 20,000 bytes, roughly 30 times today's. A real chain is measured, not assumed (Activity 6.3). 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

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

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, first published February 2026).

The Web PKI is heading toward MTCs. Checked in June 2026 and again for this edition: public root CAs have issued no ML-DSA roots that chain into the major browser and operating-system trust stores. The CA/Browser Forum has no merged requirement permitting ML-DSA in publicly trusted TLS certificates (a ballot to allow it is in draft). The fork therefore stays as written: public-facing certificates track MTC, and private PKI proceeds with X.509 PQC now. 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 are expected to follow the MTC direction that the major browser and public-CA operators have announced, on the operators' dates and subject to their trust-program decisions. 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 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

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

pays for itself in operational efficiency. The PQC readiness it enables is an additional benefit, not the sole justification.

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 raw for ECDSA, a few bytes more as DER-encoded in a certificate. 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:

- **Shorten key lifetimes immediately** – Reduce root CA key lifetimes from 20+ years to 10 years, intermediate CAs to 5 years, and end-entity certificates to 90–365 days. This limits TNFL exposure and builds organizational muscle for increased rotation frequency.
- **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.
- **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.

- **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.

The Chrome Root Program's client-authentication separation forces a trust-arrangement decision. Make it agile. The public CAs treat the Chrome Root Program Policy (v1.8) as binding. It sets two dates:

- **From June 15, 2026.** Subordinate CAs disclosed to the CCADB assert the serverAuth extended key usage only.
- **From March 15, 2027.** Publicly trusted TLS subscriber certificates must assert serverAuth only.

Client authentication therefore leaves the Chrome-trusted public hierarchies. Other root programs and private hierarchies set their own rules (Chrome Root Program Policy v1.8, section 1.3.2).

Some mTLS estates authenticate clients with publicly trusted certificates: partner APIs, service-to-service links, and VPN and 802.1X device identity on public issuers. Every one of them needs an alternative client-authentication trust arrangement inside the program's Year 1 to 2. The deadline is the root program's, not the program's own.

Three arrangements work:

- a private hierarchy the organization operates
- a managed private CA
- a separate purpose-built hierarchy

All three are called "the private hierarchy" below.

That build is the cheapest moment to make the hierarchy agile. Four things to do while building it:

- Issue from a CA whose signature algorithm is a configuration choice.
- Provision an ML-DSA-capable issuer alongside the classical one from the first ceremony (RFC 9881 profiles).
- Keep the trust-store distribution mechanism updatable in the field.
- Record the hierarchy in the CBOM as Track B assets.

Whether the relying parties in scope can already validate ML-DSA chains is a per-product question. Answer it from release notes and a test, not from a roadmap. A private PKI built in 2026 with a hard-coded classical issuer is a second migration scheduled for the moment the first one closes.

Signature placement in mixed chains. A private chain does not go post-quantum in one step. Three algorithms are in play at each element: the subject key algorithm, the issuer's signature algorithm on that certificate and, at the leaf, the CertificateVerify algorithm in the handshake.

Each placement puts something different on the wire and asks something different of the relying parties:

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

Placement	What it puts on the wire	What it needs
Post-quantum root	One post-quantum signature, the root's signature on the intermediate. The root's own public key is not sent.	Every relying party's trust store has to contain the new root before anything beneath it validates. Verifier updatability (Phase 3, Dimension 2) therefore gates this placement.
Post-quantum intermediate, signed by the existing classical root	The intermediate's public key and its signature, on each leaf.	No new trust anchor. This is the usual first move for an internal PKI, because pilots can exercise the post-quantum verification path without re-keying every leaf.
Post-quantum leaf	The leaf's post-quantum public key, and a post-quantum CertificateVerify signature, on every handshake.	No trust-store change.

The chain's integrity still depends on the classical root until the root moves, and no chain is quantum-resistant end to end while any signature edge is classical. Record the placement decision per hierarchy with the signature edges that remain classical, the measured byte and CPU budget against Activity 6.3, and the verifier inventory from Phase 1. The 2026 experimental study of signature placement in TLS 1.3 hierarchies (Delgado, arXiv 2604.06100, April 2026) measured the same point. Cost follows where a signature family appears in the chain, not whether it appears at all. SLH-DSA in the server leaf moved handshake latency and server-side compute by orders of magnitude. SLH-DSA confined to the upper trust layers, with ML-DSA in the leaf, stayed within an operable range.

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

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

Action	Timeline	Consideration
Upgrade HSM firmware to PQC-capable versions	As available	Current firmware generations from the major HSM vendors include initial ML-KEM and ML-DSA support; for some product lines PQC arrives as a new hardware variant rather than a firmware upgrade, so budget for replacement where the installed model is excluded. Model and firmware references are listed in Tool and Vendor References with an as-of date
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

The HSM and KMS register. The register records, for each module: the visible platform (the product as procured), the supplier's canonical firmware family, the manufacturer of origin, and the FIPS 140-3 or PCI HSM certificate reference. Two products under two names can be one firmware family, and a cloud KMS is a dependency boundary rather than a manufacturer. The four fields are therefore recorded separately, and the boundary-verification rule in Activity 7.7 applies to the cloud rows. Organizations using the Applied Quantum CBOM Profile record the same fields as `custody:*` properties. The register works without them.

Key Generation and Entropy

Every PQC key the estate generates, and every hybrid shared secret it derives, is as unpredictable as the entropy source behind the generator and no more. The register above assumes the following requirement.

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

- HSM and KMS key generation, and any software module that generates long-lived PQC keys, must draw from an entropy source validated under SP 800-90B. Record that validation in the register. For a current module it is an Entropy Source Validation (ESV) certificate reference. For an older module it is the ENT designation on the module certificate.
- An ESV certificate validates a source in the operating environments stated on it. It is evidence for those environments and not deployment proof: a module operated outside its tested envelope (temperature, voltage, virtualization layer) makes an entropy claim no certificate supports. Record the operating envelope with the certificate reference.
- The CMVP requirement noted in the Deployment Environment Classification applies to new procurement: IG D.K Resolution 23 (April 9, 2026) strongly recommends through December 31, 2026, and requires afterward, that physical noise sources be evaluated at multiple operating conditions. A module whose entropy validation predates the requirement stays valid. A new procurement should ask whether the source was evaluated to it.
- A "quantum" label on a randomness product is not a requirement of this framework or of any standard it cites. The buyer's questions are the validated source, its ESV or ENT reference, and the operating envelope. The requirement is the SP 800-90B validation, not the physics: a quantum random number generator without one is unvalidated. A conventional generator with one qualifies.
- In the TLS hybrid construction, the server encapsulates and the client decapsulates. The decapsulating client recovers the encapsulation randomness exactly. The server's generator output is therefore exposed to the client, in a way that an ECDH ephemeral scalar is not (RFC 10024, Security Considerations).

If that generator is insecure, an attacker can use the disclosed output to compromise its state, and with it any other secret produced from the same state. Implementations follow the random-bit-generator guidance in FIPS 203 and in RFC 9846 Appendix C.1. The register records the generator per module and per role.

Key generation is CCF layer 6 (key generation) in the Cryptographic Concentration Framework, for institutions that measure concentration there. Organizations that keep the register in a Profile-conformant CBOM record the same facts as `entropy:*` properties. The register works without either.

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

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

- Hybrid TLS 1.3 (X25519 + ML-KEM-768): ~3.5–4.5 KB (key exchange adds 1,184 bytes to the ClientHello and 1,088 bytes to the ServerHello, 2,272 bytes combined)
- 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; TLS certificate compression reduces the server's certificate message and does nothing for the ClientHello, whose hybrid key share stays at the size the group sets
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	RFC 7383 fragments only encrypted IKE messages and cannot fragment IKE_SA_INIT; RFC 9370 carries the additional ML-KEM exchange in IKE_INTERMEDIATE, which is encrypted and can be fragmented; specify the transform identifiers and the product's profile, then test fragmentation, NAT traversal and failure behavior 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

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

Protocol	PQC Challenge	Mitigation
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:

- **Baseline measurement:** Capture current performance metrics (latency p50/p95/p99, CPU utilization, memory, throughput) under representative production load
- **Lab validation:** Deploy hybrid/PQC in isolated test environment; measure same metrics
- **Canary deployment:** Deploy to 1–5% of production traffic with A/B comparison
- **SLO evaluation:** Compare canary metrics against defined SLOs from Phase 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 about 2.3 KB of key shares across both directions; PQC certificates add 2–10 KB per	Significant for high-volume TLS termination points (CDN edges, load balancers); calculate aggregate

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

Resource	Expected Impact	Planning Action
	connection	bandwidth
Storage	Larger certificates and keys in certificate stores; larger OCSP responses	Manageable; plan for 2–5x increase in certificate storage

APPLYING THIS PHASE

- **Who runs it.** Three workstreams split the work. The PKI & Code Signing workstream owns the PKI and HSM work. The Network & TLS/VPN workstream owns the middlebox and capacity work. The Embedded/IoT/OT workstream owns gateway placement. The QRPM is accountable for the schedule, and the Executive Sponsor for the capital decisions in it.
- **Minimum viable version for a small estate.** Keep the HSM and KMS register at its four fields, with certificate references. Reduce and automate end-entity certificate lifetimes. Run one middlebox test on the production edge path, using hybrid handshakes and a post-quantum certificate chain. Build the private hierarchy that the client-authentication sunset forces, issued from an algorithm-configurable CA. The performance program at this size is the pilot's before-and-after measurement, not a load laboratory.
- **The decision this phase produces.** Which infrastructure components are upgraded, replaced or bridged, and in which refresh window, taken by the SteerCo on the roadmap (Activity 4.3). An HSM procurement order that needs capital is approved by the Executive Sponsor. A component that cannot be upgraded or replaced by the roadmap's date is handled under Activity 7.4 as a vendor-as-blocker case, with the bridging pattern named.
- **What it hands on.** Phase 6 hands three things to later work:
  - The PKI modernization plan and the HSM schedule, to Phase 5, where each delivery unblocks a set of wave targets.
  - The infrastructure-layer CBOM entries, including the register and its entropy references, to the Cryptographic Record.
  - The vendor gaps found in testing, to Phase 7, with the test evidence attached.

OUTPUTS

Output	Quality Criteria

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

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

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. Each pilot exposes an infrastructure gap. Each gap closed widens the scope of the pilots that follow. These two phases form a continuous feedback loop.

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

- **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.

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 are deployed 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

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

aggregate bandwidth increase from hybrid handshakes. Model the capacity impact before production deployment.

Common failure: a validated module reported as post-quantum capability when the certificate covers a module version or an operating environment the estate is not running, or covers PQC only as a non-approved service.

CHALLENGE QUESTIONS

- Which HSMs and KMS instances in the register run a PQC-capable firmware today, which have a dated order, and which have neither?
- Which production middleboxes have been tested with a post-quantum certificate chain, and where is the result for each path?
- Which private hierarchies were issued this year from a CA whose signature algorithm is a configuration choice, and which were hard-coded classical?
- For each module certificate cited as validation, does it name the module version and operating environment the estate runs?
- What is the aggregate bandwidth and CPU increase modeled for the highest-volume termination point, and who signed off the capacity plan?

MATURITY INDICATORS

Level	Indicator
Level 0 – Unaware	No awareness of PQC infrastructure impact; no testing
Level 1 – Aware	Awareness of PKI/HSM/network challenges; no concrete plans
Level 2 – Initiated	HSMs inventoried with PQC status; PKI modernization plan drafted; initial middlebox testing underway
Level 3 – Progressing	HSM upgrades in progress; PKI dual-stack operational; all production middleboxes tested; performance baselines established for Tier-1 systems
Level 4 – Advanced	Infrastructure fully PQC-capable across IT estate (a readiness state; verified use is counted by the Coverage and Trust KPIs, not here); capacity planning validated at

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

Level	Indicator
	production scale; PKI automated with shortened lifetimes; middlebox monitoring integrated into continuous discovery
Level 5 – Optimized	Infrastructure PQC-capable across IT and OT boundaries; certificate lifecycle automated at short lifetimes; HSM, PKI and middlebox posture monitored continuously and refreshed without program effort

© 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 are 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.

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

Alternative sourcing removes a dependency only if the alternative does not share it. Whether two vendors' products sit on one implementation lineage is the question the Cryptographic Concentration Framework answers, and the vendor register carries its result by cross-reference rather than by recomputing it here.

7.2 Execute Vendor Engagement

For each strategic vendor, obtain:

- **PQC feature matrix:** Which products support which PQC algorithms; current vs. planned
- **GA version and dates:** When PQC support will be generally available (not beta, not "coming soon")
- **Interoperability guidance:** Tested configurations with common counterparties

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

- **Performance data:** Vendor's own benchmark data for PQC operation
- **Fallback behavior:** What happens when a PQC negotiation fails (graceful degradation to classical, or hard failure?)
- **CBOM data:** Cryptographic bill of materials for the vendor's product

Vendor PQC Readiness Questionnaire (core questions):

The first two questions are the agility questions. The third grades the firmness of every date. Ask these three first, because their answers decide whether the vendor's roadmap dates matter at all. A product whose algorithms change only with a new release, or with new hardware, turns every future algorithm change into a procurement.

- Can cryptographic algorithms be changed via configuration without recompilation, reinstallation, or hardware replacement, and which algorithms, parameter sets and hybrid combinations can the current release be configured to?
- Can the product's verifiers and trust stores be updated in the field (new trust anchors, new signature algorithms) without replacing the device or the installation, and by what mechanism?
- For every roadmap date you provide: is it committed (contractual, with remedies), forecast (a planning date stated in writing), or announced (a public roadmap statement)?
- Which NIST-standardized PQC algorithms (ML-KEM, ML-DSA, SLH-DSA) does your product currently support?
- For algorithms not yet supported: what is the GA date and version?
- Does your product support hybrid/composite operation (classical + PQC simultaneously)?
- What is the performance impact of PQC operation (provide benchmark data)?
- What testing have you done for interoperability with common counterparties?
- Will you provide a CycloneDX CBOM for the product? It should record the eight fields of the Activity 2.1 record that describe the product itself, and the `dependencies` relationships between them. The eight fields are:
    - component identifiers
    - algorithm OIDs or Cryptography Registry identifiers
    - key sizes and security parameters
    - protocol contexts
    - implementations with versions
    - certificate references
    - cryptographic functions
    - quantum vulnerability status

For HSM, KMS and cloud key-management products, it should also record the register fields of Activity 6.2: platform, firmware family, manufacturer of origin, and certificate reference.

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

The fields the organization records for itself (confidentiality horizon, regulatory deadline, data classification, migration status, owner, vendor dependency flag and the five fields recorded at Phase 3) are not asked of the vendor. Where the vendor uses the Applied Quantum CBOM Profile (v1.0-RC or later), the same record supplied as Profile properties is accepted. It is not required.

- What is your commitment timeline for FIPS 140-3 validated PQC implementations?
- Does the product implement any public-key algorithm not standardized by NIST or an equivalent national authority, and under what published analysis?

For certificate authorities, four further questions, because an issuer that cannot issue on the dates the roadmap needs gates every Track B migration behind it (Activity 3.3):

- On what dates will you issue end-entity certificates signed with ML-DSA, and with FN-DSA?
- Do you support, or when will you support, hybrid or composite certificate issuance?
- What is your roadmap for Merkle Tree Certificates for the public Web PKI (Phase 6, PKI Architecture Fork)?
- By what date will your own signing keys be PQC-capable, and how will the transition of your roots and intermediates be communicated to relying parties?

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
- Required: CBOM handling terms, covering how the supplier stores, transmits and retains the CBOMs it receives from the customer and produces for its products, with access limited to named roles and the retention period stated (Activity 2.5)

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

- Required: no public-key algorithm outside the NIST standards, or the standards of an equivalent national authority, in approved-mode use. Any other public-key algorithm in the product is disclosed with its published analysis

Crypto-agility as a contractual requirement with an acceptance test. The clause above already warrants configuration-driven algorithm change. It must be tested rather than warranted: acceptance testing before go-live includes a rehearsed algorithm change on the product, executed by the customer's team using the rehearsal pattern in Activity 5.2, with the elapsed times recorded and the rollback exercised. A product that passes the functional acceptance tests and fails the rehearsal has not met the clause.

Supervisory outsourcing guidance in some jurisdictions has begun to require contractual crypto-agility in the same terms, and the acceptance test is what turns that clause into evidence.

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. CISA's guidance on procurement-relevant product categories signals where contract expectations are heading. Legal teams move faster when handed a government-grade exemplar.

No non-standardized public-key algorithms. Procurement excludes from approved-mode use any public-key algorithm that is not standardized by NIST or by an equivalent national authority, and the vendor questionnaire (Activity 7.2) asks the question directly. A product that uses a proprietary or pre-standard algorithm as its "quantum-safe" claim has added an unreviewed dependency to the estate. There is one caveat, and it is sovereignty. A national algorithm mandated under a documented regime counts as standardized by an equivalent authority for the purposes of this rule. The Algorithm Sovereignty note in the Regulatory & Standards Alignment Map explains which regimes qualify. Record the algorithm's use together with the regime that requires it. The verification fits in an afternoon:

1. **Public specification.** The algorithm is fully specified in a document anyone can read, with test vectors.
2. **Public cryptanalysis.** It has been analyzed in the open literature by parties other than its designers, for long enough that the analysis has had time to fail.
3. **Standards-body track.** It is standardized, or on a published standardization track, by NIST or an equivalent national authority.
4. **Validated implementation.** A CAVP-tested implementation exists, or is on a dated track, and a CMVP-validated one where the deployment environment requires it.

A product that fails the first two tests is excluded. The third and fourth set the date on which it may enter approved-mode use.

© 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:

- **Deploy bridging patterns:** Software PQC overlay (gateway TLS termination, PQC key-wrapping around classical HSM, double-encryption layer) until native vendor support arrives.
- **Run champion-challenger:** Evaluate an alternative vendor through POC while maintaining the incumbent. Communicate the POC to the incumbent vendor as leverage. A challenger that shares the incumbent's implementation lineage removes no dependency. Check the lineage before the POC (Cryptographic Concentration Framework).
- **Escalate contractually:** Invoke SLA remedies, maintenance credit provisions, or right-to-replace clauses.
- **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

Firmness grading. Every roadmap date recorded in the vendor scorecard carries one of four grades: committed (a contractual or written commitment with remedies), forecast (the vendor's planning date, stated in writing without remedies), announced (a public roadmap statement) or unknown (no date obtained). Phase 3 Dimension 3 scores the firmness grade, not the calendar date, because an ungraded date implies more certainty than exists.

The scale aligns with the `vendorReadinessStatus` values of the Applied Quantum CBOM Profile for organizations that record vendor data in a Profile-conformant CBOM.

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

Institutions running the Cryptographic Concentration Framework use its Instrument Lineage Disclosure Request as the companion for implementation ancestry. This framework does not include it.

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)

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. Six elements make up the pattern.

- **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.

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

7.7 Cloud Shared Responsibility and the SaaS Class

Cloud has appeared in this framework mainly as tooling surface (CSPM queries, KMS checks). For migration, the sharper question is 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 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.

Most programs miss 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.

Reading a provider's roadmap. Ask every cloud provider for both halves of the disclosure: the current negotiated state per service (which key-exchange groups and signature algorithms each managed endpoint offers today, observable at the boundary) and the dated targets per service, graded for firmness (Activity 7.5). A provider's completion date is the date the provider's half of the boundary migrates. It is not the customer's date, because everything above the line in the previous paragraph is untouched by it.

Read the exclusions in the footnotes of every roadmap before marking any provider-side component migrated: regions, service tiers, legacy endpoints and customer-managed key options are the usual carve-outs, and a component any of them covers is unmigrated whatever the headline says. The two-part ask goes into the Activity 7.2 questionnaire for every cloud and SaaS vendor.

SaaS is a control-posture class of its own, not a variant of cloud. There is usually no bridging pattern for a SaaS application, because the customer cannot terminate its TLS, wrap its keys or proxy its internals.

Three exceptions exist: client-side encryption where the service supports it, customer-managed keys, and a customer-controlled endpoint. Each of these changes what the customer controls. None of them changes anything on the provider's side. 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.

APPLYING THIS PHASE

- **Who runs it.** The Vendor Orchestration workstream runs the engagement and the scorecard. The Policy/Compliance/Procurement workstream owns the clauses and the RFP template. Three people are accountable: the QRPM for the register, the procurement lead for clause insertion, and the commercial relationship owner for each counterparty decision under Activity 7.6.
- **Minimum viable version for a small estate.** Send the questionnaire to the top ten vendors from Phase 0 scoping, and include the two agility questions, the firmness question and the certificate-authority questions. Insert the three mandatory clauses into the standard contract template. Pre-draft one bridging pattern, for the single vendor most likely to block. Keep the scorecard as a spreadsheet with firmness grades, updated quarterly.
- **The decision this phase produces.** The vendor classification, and for each Strategic Blocking vendor that misses a date, the choice among bridge, replace and accept, taken by the SteerCo with the risk owner. The refusal decision for a counterparty, taken by the relationship owner on the program's evidence.
- **What it hands on.** Phase 7 hands four things to later work:
  - The vendor register with firmness grades, to Phase 3 (Dimension 3) and to Phase 4 (roadmap dependencies and re-baseline triggers).
  - Each GA delivery, to Phase 5, as an unblocked set of wave targets.
  - HSM, PKI and network vendor data, to Phase 6.
  - Every CBOM received under the questionnaire, to the Cryptographic Record, with the handling terms of Activity 2.5 applied.

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

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

Output	Quality Criteria
Vendor governance cadence	Operating rhythm established with monthly/quarterly reviews

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 in parallel with

Every phase, 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

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

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.

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.

Common failure: a roadmap date recorded in the register as a deployed state, and reported upward as one.

CHALLENGE QUESTIONS

- Which vendor dates in the scorecard are committed, with remedies, and which are forecast, announced or unknown?

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

- For each Strategic Blocking vendor, which bridging pattern is deployed or pre-drafted, and what is its expiry date?
- Which vendors answered the agility questions with configuration-driven algorithm change, and which answered with a new release or new hardware?
- Which provider-side components are marked migrated on observed negotiation at the boundary, and which on a published roadmap?
- Which CBOMs received this quarter contain the minimum record of Activity 2.1, and which arrived as a spreadsheet?
- Where a vendor's firmware family is the same implementation as the alternatives, what does the Cryptographic Concentration Framework say about substitution?

MATURITY INDICATORS

Level	Indicator
Level 0 – Unaware	No vendor engagement on PQC; assumption that "vendors will sort it out"
Level 1 – Aware	Ad-hoc inquiries to a few vendors; no structured tracking
Level 2 – Initiated	Top 10 vendors formally engaged; questionnaires sent; responses tracked; vendor criticality classification complete
Level 3 – Progressing	PQC in standard procurement language; contracts include dated commitments and remedies; bridging patterns deployed for blocked systems; vendor scorecard reported to SteerCo
Level 4 – Advanced	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
Level 5 – Optimized	Vendor governance is a permanent BAU function; bridging patterns eliminated; every strategic vendor's delivery verified against committed dates; implementation lineage checked before any substitution

© 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 EVIDENCE STANDARD

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:

- **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. Passive observation shows the negotiated key-exchange group, which is sent in the clear in the ClientHello and ServerHello. In TLS 1.3 it does not show the certificate chain or the signature algorithm, which are encrypted, and a cipher-suite value identifies neither. Track B evidence therefore comes from endpoint telemetry, the certificate inventory and authorized active tests.
- **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.
- **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.
- **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.
- **Agility rehearsed, not asserted.** For Tier-1 systems, the evidence includes a rehearsal record showing that the algorithm was changed by configuration in staging, verified by observed negotiation, and rolled back, with the elapsed time per

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

segment (Activity 5.2). A system migrated without this record is migrated once. The record is what shows it can be migrated again on a planned date.

- **Incidental capability counted separately.** Capability that arrived through a package default (hybrid key exchange enabled by a library update) is recorded in the CBOM as hybrid deployed and reported on a separate "incidental" line until the negotiated outcome is observed and the policy is enforced and tested. A package manager's defaults never read as program progress.
- **Hybrid negotiated, fresh or resumed; never inherited.** For long-lived channels, the verification record shows a hybrid group negotiated after hybrid enablement, in a fresh handshake, or in a resumed one that included a fresh hybrid key share, and shows that PSK-only resumption and 0-RTT are disabled for the flows in scope. A connection already open on a hybrid-capable endpoint before the change is unmigrated until it re-handshakes (Activity 5.2).
- **No draft group negotiated.** The negative test of Activity 5.2 is repeated at verification: no pre-standard code point (the Kyber768 draft groups obsoleted by RFC 10024) is offered or negotiated on any verified endpoint.
- **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.

Evidence by Cryptographic Function

The evidence standard above is written for negotiated protocols, where the handshake provides the evidence. The matrix below states, for each cryptographic function in the record (the cryptographic-function field of Activity 2.1), what counts as positive evidence of migration, which negative or boundary test shows the classical path is closed, and which residual exposure is recorded even when both are met. Activity 4.6 reads the first two columns as the first two conditions of the agility criterion for every function that is not a negotiated protocol, and Activity 5.4 reads them at wave exit. The residual-exposure column is entered in the verification record and reported with the protection result. It is never subtracted from a coverage figure.

Function	Positive evidence	Negative or boundary test	Residual exposure to record
Key establishment (TLS, IKE, SSH)	The hybrid or PQC group observed negotiated in production traffic, fresh or resumed with a fresh hybrid key share (Activity 5.2)	A classical-only peer is refused where policy requires it; no pre-standard code point is offered or negotiated	Sessions recorded before the change; peers still on classical fallback where policy permits it, listed by counterparty

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

Function	Positive evidence	Negative or boundary test	Residual exposure to record
Certificate-based authentication	Production relying parties validate the PQC or composite chain end to end, shown by endpoint telemetry and the certificate inventory, with revocation exercised	A classical-only chain is rejected by the relying parties in scope where policy requires it; a relying party that cannot validate the chain is listed, not assumed	Relying parties outside the organization's control that still trust the classical chain; the retired chain wherever it remains in a trust store (Decommissioning Classical Material)
Firmware, update and secure boot	A post-quantum-signed image (SP 800-208 or ML-DSA) accepted by the production verifier population, with the signer's state management evidenced where the scheme is stateful (SP 800-208 section 8)	An image carrying only a classical signature is rejected by a sampled verifier; where dual signatures are in use, the verifier policy requires both	Fielded verifiers that cannot be updated and still accept the classical signature, counted against the verifier population; the classical signing key wherever it stays trusted
Long-term document and evidence preservation	Signatures, timestamps or seals applied or re-applied under a PQC scheme, with a validation record under the applicable profile (Appendix H; GRC Implementation, Qualified Trust Services and Electronic Signature Law)	A record carrying only a classical signature fails validation under the policy that governs new signatures, or is archived with a dated attestation that carries it into the PQC regime	Records that cannot be re-signed and rely on preserved classical validation material, with the period over which they remain relied upon
Stored data and wrapped keys	Data or key material re-wrapped or re-encrypted under a PQC-derived or	Decryption with the retired wrapping key fails, or the retired key is shown	Copies taken before the change are not repaired by re-encryption; the

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

Function	Positive evidence	Negative or boundary test	Residual exposure to record
	symmetric key of adequate strength, with the disposition of the previous wrapping key recorded (Activity 5.6)	destroyed or held decrypt-only under controlled custody	period and the data exposed are recorded as exposure that persists
Application and account authorization	Tokens, assertions or credentials issued and accepted under a PQC or composite signature (COSE and JOSE identifiers, certificate signature algorithms), verified at the accepting service	A token carrying only a classical signature is rejected by the accepting service where policy requires it, tested at the service and not inferred from the issuer's configuration	Sessions and refresh tokens issued before the change and still valid; accepting services outside the organization's control
Proof systems and commitments	The proof system, commitment scheme or signature scheme in use rests on no quantum-vulnerable assumption, with the setup it runs under recorded	A proof or commitment produced under the retired scheme is not accepted by the verifying party; where the verifier is a protocol the organization does not control, the acceptance rule is recorded as dependent (Activity 7.6)	Commitments and proofs already published under the retired scheme, and any trusted setup whose exposure is irreversible, with the delivery state of the protocol change that closes them

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 closes with an approved disposition for every superseded key, certificate and trust relationship, and disposition is not always destruction. Decryption keys stay under controlled, decrypt-only custody for as long as retained archives need them (Activity 5.6). Public verification material and historical validation evidence are preserved where signed records must remain verifiable (Appendix H), and the classical component of an approved hybrid deployment stays active by design.

For each completed migration, do five things:

1. Revoke superseded certificates on schedule, rather than letting them expire silently.
2. Destroy retired private keys under the organization's key destruction standard. Obtain destruction certificates for HSM-held keys.
3. Remove retired roots and signing certificates from trust stores. A key is only retired when nothing trusts it.
4. Remove classical-only cipher suites from configurations, once counterparty support permits.
5. 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.

PROGRAM CLOSURE AND TRANSITION TO BAU

Define closure criteria in advance. There are three:

1. Maturity Level 4 or higher across all seven domains.
2. A verification disposition for every Tier-1 and Tier-2 service. A disposition is one of four states: migrated; migrated with agility debt and a funded closure date; enclosed or compensated under an accepted exception; or retired.
3. The closure proof. The closure proof is a rehearsed second algorithm change on at least one Tier-1 system per track, executed inside the agility window (the 48-hour assessment plus the change window the organization has defined), minuted, with the segment timings filed in the evidence dossier, and the agility-debt register at zero unaccepted debt.

Every remaining entry has either a funded treatment with a date or a SteerCo-signed exception. Without the closure proof the program has completed a migration and has not demonstrated the capability it was chartered to build.

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

Zero unaccepted debt is a governance condition, not debt elimination. The register at closure can still record substantial debt that is funded or excepted. The condition says only that none of it is unowned.

The closure report therefore shows five things, on separate lines that the closure decision cites:

1. **The services that remain exposed on a classical path.** List the unresolved exposed services by tier. For each one, give the exposure recorded in the residual-exposure column of the verification matrix.
2. **The compensating controls in place for them.** Verify each control by observed negotiation, or by the function's boundary test. A control's presence on a change ticket is not verification.
3. **The residual risk the accepting body has accepted.** Name the accepting body.
4. **The remediation date for every entry.**
5. **The agility-debt total.** Count funded entries and open entries in full. Only the unaccepted count has to be zero.

A verification disposition on every Tier-1 and Tier-2 service is a completeness statement about the record. It shows that every service was examined and placed in one of the four states. It is not a claim that every service migrated. Read the migrated count from the first state alone.

An internal exception can accept internal risk and nothing more. An exposed service may also be subject to an external obligation, such as a regulatory deadline, a scheme rule or a contractual term. The accepting body cannot waive an obligation it has no authority to waive. The closure report marks that entry as an open obligation with its external date. It is never marked as accepted residual risk.

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 Crypto-Agility workstream transfers as a continuing capability with a funded owner and a standing rehearsal cadence, and the Cryptographic Record owner named in Phase 0 keeps the CBOM.

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 proof	Rehearsed second algorithm change on at least one Tier-1 system per track inside the agility window; minuted; segment timings in the evidence dossier; agility-debt register at zero unaccepted debt (every remaining entry funded with a date or under a SteerCo-signed exception)
Closure report	Unresolved exposed services by tier with their residual exposure; compensating controls verified in operation; accepted residual risk with the accepting body named; remediation dates for every entry; the agility-debt total with funded, open and unaccepted counts; external obligations marked open with their dates, never accepted internally
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

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. The first auditor, regulator, or counterparty who asks for proof will treat it accordingly.

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

- **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.

THE SEVEN PROGRAM FOUNDATIONS

The eight phases above describe *what* the program does and in *what order*. Several capabilities belong to no single phase and operate continuously across the entire program lifecycle. A program that executes the phases without building these capabilities will produce a migration that cannot be measured, sustained, or staffed, and that will not satisfy regulators.

This section covers seven cross-cutting 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, the skills and team structures needed to execute the work, SOC implementation and GRC implementation. Each should be established early, ideally during Phase 0 and Phase 1, and maintained throughout the program's lifetime.

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

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

Level	Name	Description
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 and verified in use; PKI modernized, which is a readiness state reported apart from the verified Trust count; 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
Pilots & Deployment	No pilots	Lab testing only	2+ production pilots with	Tier-1 systems on	Estate-wide deployment,

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

Domain	Level 1	Level 2	Level 3	Level 4	Level 5
			measured results	hybrid/PQC, verified in use; wave rollout underway	verified in use; 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 change rehearsed on at least one Tier-1 system	Algorithm swap by configuration for Tier-1 systems, rehearsed within the agility window; automated cert lifecycle	Organization-wide agility; algorithm changes are routine operations with a standing rehearsal cadence
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

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

Self-Assessment Scoring

For each domain, score 0–5; Level 0 in every domain is Unaware, as in the Maturity Levels table. The Maturity Indicators tables at the end of each phase use the same six levels and names. Overall maturity level is the lowest individual domain score (the weakest domain constrains overall readiness), and the closure criterion reads that scale.

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

Re-assess semi-annually. Report the seven-domain profile and its trend as well as the composite, which is the lowest domain score and by itself hides which domain is holding it. A domain that is flat across two consecutive assessments is an escalation trigger to the SteerCo, whatever the composite shows.

METRICS, KPIS & REPORTING

Board-Level KPI Pack (Report Quarterly)

KPI	What It Measures	Target
Coverage	% of Tier-1 endpoints, internet-facing and internal, using hybrid/PQC key exchange under the enforcement counting rule below	Year 1: 10%; Year 2: 60% (the Year-2 roadmap milestone covers the internet-facing Tier-1 endpoints; internal Tier-1 endpoints follow in Year 3, Activity 4.2); Year 3: 95%
Trust	# of PKI/signing anchors whose post-quantum signatures are verified in use at their relying parties (the certificate-based authentication and the firmware, update and secure boot rows of the verification matrix); anchors that are PQC-capable but not yet verified in use are reported on a separate readiness line,	Year 1: every root and signing anchor inventoried with its verifier updatability recorded; Year 2: a post-quantum issuer provisioned in each private hierarchy and dual-signed code or firmware accepted by a sampled production verifier; Year 3: Tier-1 anchors verified in use at their relying parties. Lifecycle line: Year 1 root CAs

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

KPI	What It Measures	Target
	and lifetime reduction and automation on a separate lifecycle line, because neither a capability nor a shorter lifetime changes the signature algorithm a relying party accepts	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%

The five KPIs are the board's view of progress. The board-level KRIs in GRC Implementation are the risk view beneath them. The quarterly board paper carries the KPI pack, with KRI breaches reported as exceptions.

Risk-weighted coverage (companion metric). The Coverage KPI counts endpoints equally, which makes 60% read as progress even when the highest-value flows are all in the unprotected 40%. Report an exposure-weighted companion over the same population, the same quarterly window and the same enforcement counting rule as the Coverage KPI. The unit is the weighted endpoint: weight each endpoint by its Dimension 1 score, which reads the confidentiality need recorded under Activity 1.1 and never the retention obligation, on the ratified Low = 1 to Critical = 4 mapping (Activity 3.2).

The figure is the summed weight of the covered endpoints over the summed weight of the population. The weighting is published with the figure. An endpoint is covered only when it meets the enforcement counting rule, so an endpoint with partial protection contributes nothing to the numerator. An endpoint whose Dimension 1 inputs are unknown is unscored under the rule on unknown exposure in Activity 3.1: it enters neither sum and is reported on its own line with its count, never defaulted to the lowest band.

The inputs already exist, in the Phase 1 classification and the Phase 3 scoring run. Offer the companion alongside the plain endpoint metric, not as a replacement. The pair

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

answers the question a regulator will actually ask, which is not how many systems, but how much of what matters.

- **Enforcement counting rule.** A system counts toward Coverage when the negotiated outcome is observed in production traffic and the policy is enforced and tested (classical-only refused where policy requires it). Capability that arrived through a package default is recorded in the CBOM as hybrid deployed and reported on a separate "incidental" line until the observation and the enforcement test exist (Migration Verification & Program Closure).
- **Denominator changes.** When improved discovery grows the denominator of any coverage metric, the change is reported as its own line with its cause, and the prior figure is restated on the new denominator. A coverage drop caused by finding more of the estate is a discovery success and is never presented as migration regression. The Cryptographic Concentration Framework's treatment of movement between assessment cycles (C.8) is the pattern.
- **Endpoint coverage and path coverage.** The Coverage KPI and its risk-weighted companion count the same endpoints, the companion weighted as above, derived from the rows of the record as Activity 2.1 defines them (the rows that share one component identifier and one network-facing protocol context). The Phase 3 scored-entry count counts the rows themselves, and both figures are read from the same population. Whether a service's traffic is protected end to end across every path and counterparty it uses is the Cryptographic Concentration Framework's Effective Coverage. Institutions running both report the pair.

Agility KPIs (CISO cascade, reported monthly; summarized to the board under the Agility KPI).

KPI	Measurement	Threshold
Configuration-driven share	% of Tier-1 systems whose key-establishment and signature algorithms change by configuration, verified by rehearsal rather than architecture review	Target by year from the roadmap's agility milestones (Activity 4.2). A year missed by more than 10 points triggers Crypto-Agility workstream review
Rehearsed time-to-change	Elapsed time across assessment, decision, change and verification on the most recent rehearsal, per track, Tier-1 systems (Activity 5.2)	Inside the agility window (the 48-hour assessment plus the defined change window). A result outside it is a finding, not a variance
Agility-debt burn-down	Entries closed on the Phase 3 agility-debt register, per	Any quarter with no reduction triggers workstream review;

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

KPI	Measurement	Threshold
	quarter	before closure every remaining entry must carry a funded treatment with a date or a SteerCo-signed exception (zero unaccepted debt)
New-system agility compliance	% of new deployments passing the CI/CD crypto-policy check and the provider-pattern review	Below 95% in any month triggers architecture governance review (the Crypto-Agility Adoption KRI in GRC Implementation)

Operational KPIs (Report Monthly to SteerCo)

KPI	Measurement
Migration tasks completed vs. plan	Earned value / planned value (by workstream)
Pilot performance delta	p95 latency increase vs. baseline (illustrative target: <2 ms for TLS, set from the Activity 6.4 baseline)
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

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

- 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 occurs, 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 evidence that a court or a regulator will read.

Its weight depends on the decisions it records, and on whether three artifact types were captured deliberately:

- Risk-acceptance decisions, with their rationale and their re-evaluation dates.
- Budget requests, including the ones that were denied. A logged, declined funding request is evidence that the risk was raised.
- The annual record of threat reassessment, showing that the organization adjusted as information evolved. A documented decision to defer can be examined on its reasoning; undocumented inaction cannot be defended at all.

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, a design heuristic. "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 law is a design heuristic: a statement of direction with no units and no calibration, measured in practice by the rehearsed time-to-change (Metrics, KPIs & Reporting), and no budget figure is derived from it.

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

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; HAWK was withdrawn in July 2026, leaving eight. HQC is being standardized as a non-lattice backup for ML-KEM. Daniel Bernstein's research has raised questions about the robustness of specific ML-DSA implementations. 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.

The same investment hedges three events, and they arrive on different clocks:

- a CRQC
- a standardized algorithm weakened under classical or AI-assisted cryptanalysis
- the routine deprecation cycle that retired SHA-1 and 3DES AI-assisted cryptanalysis shortens the interval between a mathematical idea and a practical attack, the interval that agility exists to survive. The claims the Skills-agility test below asks the organization to assess within 48 hours will increasingly be machine-generated and arrive in volume, which is a reason to build the assessment as a process rather than rely on an individual's reading time.

The Five Dimensions of Crypto-Agility

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 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?

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

- **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. Security engineers who understand the provider-pattern architecture, and a CTI function that can interpret cryptanalytic developments, meet this requirement without PhD cryptographers on staff. The test: can the organization assess the security implications of a new algorithm within 48 hours of its publication? The 48-hour assessment plus the change window the organization defines (ten business days by default, the Dimension 2 test) is the agility window referenced at the gates, in the rehearsal method and at closure.
- **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 providers (the ENGINE API is deprecated in OpenSSL 3), 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 carry both classical and post-quantum algorithms and negotiate the approved hybrid construction wherever the counterparty offers it. Dual capability, hybrid combination and fallback are three different things. Carrying both algorithm families does not mean every connection combines them. A connection that settles on classical-only is a fallback and not a hybrid. Classical-only negotiation is permitted only under an explicitly authorized exception, recorded with its authority and its expiry. The connection is excluded from the Coverage KPI under the enforcement counting rule (Metrics, KPIs & Reporting). Where protection policy requires refusal, the refusal is configured and tested rather than left to degrade.

© 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.

Mechanisms by Layer

The principles above state what an agile estate looks like. The catalogue below states what the Crypto-Agility workstream does to get there and how each mechanism is verified, so that funding and audit run against a list. Every verification is a test that a team can run, not a review that a team can pass by assertion.

Mechanism	What to do	How to verify
Modularity	Route every cryptographic call through an internal API or a crypto service; no direct algorithm calls in application code (Principle 1)	A build that swaps the provider passes the test suite with no application change
Policy and mechanism separated	Keep algorithm policy (approved groups, signature algorithms, parameter sets) in machine-consumable configuration profiles that services consume at start-up or on reload (Principle 2)	A policy change rolls out to a Tier-1 service with no code change and no rebuild, and observed negotiation confirms it
Hybrid and composite as the transition form	Deploy hybrid key exchange and composite or dual signatures so that the classical and post-quantum components can be changed independently (Activity 5.2)	Each component can be removed or replaced by policy; a negative test shows the retired component is no longer negotiated

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

Mechanism	What to do	How to verify
Algorithm identifiers and negotiation everywhere	Use protocols and formats that include algorithm identifiers and negotiate (TLS groups and signature schemes, IKEv2 transforms, COSE and JOSE identifiers, certificate signature algorithms); no custom format with an implied algorithm	A peer offering only a retired algorithm is refused where policy requires refusal; a new identifier is accepted after a policy change alone
Key-management abstraction	Expose algorithm and key type as KMS and HSM parameters; applications reference keys by name and purpose, never by algorithm	A key of a new type is created and used by an application through the existing interface; a failover to a backup HSM partition succeeds on its scheduled exercise, recorded as a failover and state-management result and not as an algorithm change
Certificate lifecycle automation	Automate issuance, renewal, rotation and revocation at the velocity the public-certificate lifetime schedule requires from March 2029, 47 days (Phase 6)	Mean time to rotate an end-entity certificate under one hour with zero manual steps, across a representative sample
Gateway or crypto-service pattern	Front populations that will never be agile natively (embedded, legacy, OT) with a gateway or aggregation layer that terminates or wraps their cryptography (the OT & Critical Infrastructure Extension's Encapsulate strategy; the bridging patterns in Activity 7.4)	The gateway's algorithm changes by configuration while the population behind it is untouched; the CBOM records the population as migrated by enclosure
Algorithm tests in CI/CD	Test cryptographic configuration changes in the same pipeline as code:	A deliberate policy regression is caught by the pipeline

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

Mechanism	What to do	How to verify
	staging, canary, monitoring, rollback (Principle 5)	before production
The rehearsal	Exercise all of the above on a schedule (Activity 5.2)	Elapsed time per segment recorded and inside the agility window

The Rehearsal Method

Agility is measured by rehearsal, not by architecture review. Each rehearsal records four elapsed times: assessment (from the trigger to the decision that a change is needed; the 48-hour assessment test at 48 hours), decision (approval under the standing change authority), change (rollout to the target system) and verification (observed negotiation of the new algorithm, with the rollback exercised). The sum is the rehearsed time-to-change, reported per track for Tier-1 systems (Metrics, KPIs & Reporting), and the segment that dominates it is where to invest next.

A system whose rehearsal cannot be scheduled because the change requires a code change, a rebuild or a vendor release is agility debt (Phase 3), whatever its architecture review says.

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 end-to-end algorithm swap drill on a non-production Tier-1 replica. The per-wave staging rehearsals of Activity 5.2 begin in Year 1 and are narrower exercises. 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

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

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.

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," and updated it on June 29, 2026 as CSWP 39-upd1. The update's Appendix C records a correction to the transition discussion in section 2.3 and no change of the paper's status. The paper is guidance, not a mandate:

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

it describes strategies and practices for crypto-agility across architecture, operations, governance, risk management and supply chains, and it creates no obligation.

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
- integration of cryptographic governance into enterprise risk management

CSWP 39 also sketches an illustrative maturity approach that adapts the four tiers of the NIST Cybersecurity Framework (Partial, Risk-Informed, Repeatable, Adaptive), with the caution that no single model suits every organization. The framework's six-level (0 to 5) Crypto-Agility domain is consistent with that approach. It is not endorsed by it.

Where this framework extends beyond CSWP 39 is in four areas. First, Marin's Law states the direction of the relationship between built-in agility and migration effort, and the rehearsal method gives it a measurement, the rehearsed time-to-change, that CSWP 39's maturity approach does not. Second, the five-dimension model (architectural, operational, governance, skills, supply chain) provides a broader assessment than CSWP 39's tier-based approach.

Third, the implementation roadmap integrates agility-building into the migration program's existing phase structure rather than treating it as a separate initiative. Fourth, CSWP 39 offers its tiers as illustrative and expects organizations to adapt their own risk-management frameworks to them. The mechanisms catalogue and the rehearsal method above, with the agility KPIs and thresholds in Metrics, KPIs & Reporting, are this framework's measurable implementation of that expectation.

CSWP 39's treatment of crypto-agility as an organizational capability supports every Phase 0 business case that frames PQC migration as the first exercise of that capability. The support is guidance from a standards body and not an endorsement of this framework. The algorithm swap is only the first exercise of the capability.

REGULATORY & STANDARDS ALIGNMENT MAP

Mapping Framework Phases to Regulatory Requirements

The map states the mechanism: three anchor instruments, one generic row for the sector supervisory expectations that now recur across jurisdictions, the deadline archetypes and the enforcement tiers. The instruments and dates by jurisdiction are in the Regulatory Timeline Context table at this edition and, kept current, on the Global PQC

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

Migration Clock page updated from the PQC Migration Brief. The Regulatory Horizon Report (GRC Implementation) is where each organization records the instruments that bind it.

Instrument	Relevant Phase(s)	What It Requires	Framework Output That Satisfies It
NIST IR 8547 (Initial Public Draft, November 2024; still a draft at this edition)	All	Deprecate quantum-vulnerable public-key algorithms after 2030; disallow after 2035	Full migration execution per roadmap; the regulatory deadline field on every CBOM entry (Activity 2.1)
NSA CNSA 2.0	Phase 5, 6, 7	NSS acquisitions CNSA 2.0 compliant from January 1, 2027; equipment that cannot support CNSA 2.0 phased out by December 31, 2030; CNSA 2.0 algorithms mandated for use by December 31, 2031 (CNSSP 15, updated 2024 and published March 2025, with stated exceptions); per-category chart: software and firmware signing and networking equipment 2030; web, cloud and operating-system platforms, niche equipment, and custom applications and legacy equipment 2033, the chart's last date	Procurement requirements; vendor compliance verification; Track B SP 800-208 deployment
EU Coordinated Roadmap	Phase 0, 1, 4, 5	National strategies and first steps by end of 2026; high-risk use	Charter and roadmap; discovery outputs; wave plan

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

Instrument	Relevant Phase(s)	What It Requires	Framework Output That Satisfies It
		cases migrated by end of 2030; the remainder by end of 2035	aligned to the 2030 and 2035 dates
Sector supervisory expectations (plan, procurement gate, completion)	Phase 0, 4, 7	A documented transition plan by a date; PQC requirements in new procurement from a date; named system classes completed by a date; increasingly, readiness indices and surveys against which the plan is scored	Charter and roadmap (Phase 0 and 4); procurement clauses and vendor scorecard (Phase 7); the evidence dossier

Deadline Archetypes

Every instrument on the roster sets one of three deadline types. The archetype, not the jurisdiction, determines which framework output satisfies it, so the roadmap is planned against archetypes and the Regulatory Horizon Report assigns each binding instrument to one.

Archetype	What it requires	Anchor example	Enforcement tier
Plan-by	A documented transition plan, inventory or strategy by a date	EU coordinated roadmap: national strategies and first steps (awareness, inventories, pilots) by end of 2026	Tier 4, hardening to tier 2 or 3 through national transposition
Procure-by	New acquisitions must carry PQC support from a date	CNSA 2.0: new National Security System acquisitions CNSA 2.0-compliant from 2027	Tier 1, a procurement gate with a named enforcer

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

Archetype	What it requires	Anchor example	Enforcement tier
Migrate-by	Named system classes migrated by a date	NIST IR 8547: quantum-vulnerable public-key algorithms deprecated after 2030 and disallowed after 2035 (still a draft at this edition); EU coordinated roadmap: high-risk use cases by end of 2030, the remainder by end of 2035	Tier 4, binding where a federal instrument or a sector standard cites it

Dates concentrate between 2026 and 2031 in every roster, and a program that begins in 2027 cannot complete a 4-to-15-year migration before its first binding date has passed. The window for beginning is now. Standards-body and industry dates that constrain the roadmap without being regulatory deadlines (the FIPS 140-2 sunset, past since September 21, 2026, and the public-certificate lifetime schedule) are treated where they bind: the deployment environment classification in Phase 5 and the certificate lifetime convergence in Phase 6.

Enforcement Tiers

Instruments are tagged with the tier that describes how they bind, and Phase 3 Dimension 4 reads the tag. "Mandatory deadline" means an instrument at tier 1 to 3. "General expectation" means tier 4 or 5.

Tier	How it binds	Anchor example
1 – Market access	A certification or procurement gate: without the capability, the product cannot be sold into or bought by the regulated party	CNSA 2.0's 2027 acquisition requirement
2 – Supervisory directive with a named enforcer	A supervisor with examination and sanction powers has stated a dated expectation	Sector supervisory letters and readiness assessments (see the roster page)
3 – Binding government	Statute, executive order or directive with reporting and	The federal high-value-asset dates illustrated in the Two-

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

Tier	How it binds	Anchor example
mandate with an audit trail	audit obligations	Track Migration Model
4 – Authoritative guidance	A standards or national body's published timeline that other instruments cite	NIST IR 8547's 2030 and 2035 dates; the EU coordinated roadmap
5 – Advisory framework	Recommendations without a stated enforcer or date	Sector-body playbooks and maturity models

Guidance hardens through three channels the tiers do not capture directly: contracts (a customer writes tier-4 guidance into a tier-1 obligation for its suppliers), insurers (underwriting questions convert guidance into premium and exclusion terms) and auditors (external audit scope treats published guidance as the standard of reasonable care). The Regulatory Horizon Report tags each instrument with its tier and notes which channel is hardening it.

Regulatory Pattern Since v2.1

Since the June 2026 edition, the change in the regulatory picture is in what instruments require more than in which instruments exist, and four patterns govern planning. National mandates now institutionalize the two-track split with separate dates for key establishment and for signatures (the illustration is in the Two-Track Migration Model). Supervisors publish readiness indices and sector surveys, so the claim that no one has started is checkable against data the regulator holds.

Contractual crypto-agility has entered supervisory outsourcing guidance, so the procurement clause in Activity 7.3 is now a supervisory expectation in some sectors rather than only this framework's recommendation. And mandatory CBOM guidance and submission regimes are appearing, with the federal CBOM minimum-elements direction (Activity 2.1) the first with a date. The instruments and dates behind each pattern are listed by jurisdiction in the Regulatory Timeline Context table, and kept current on the Global PQC Migration Clock page. The three anchors named above are the ones this framework is built against.

Regulatory Developments Recorded at v2.0 and v2.1 (history)

- **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. The proposal would oblige Member States to include PQC transition policies in their national cybersecurity strategies; entity-level transition-plan requirements would

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

follow through transposition and supervisory practice. 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):** Guidance on strategies and practices for crypto-agility, updated June 29, 2026. It creates no obligation. See the Crypto-Agility section for detailed alignment.
- **NIST Additional Digital Signatures – Third Round (May 2026):** Nine candidates advanced on May 14, 2026 (FAEST, HAWK, MAYO, MQOM, QR-UOV, SDitH, SNOVA, SQIsign, UOV); HAWK was withdrawn by its team on July 29, 2026. 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, confirmed per deployment profile against the applicable authority rather than assumed to satisfy every jurisdiction:

Deploy hybrid cryptography (classical + PQC) as the default for all Tier-1 traffic. Hybrid satisfies the jurisdictions that recommend it (BSI Germany; ANSSI France, whose security visa requires it for certified products). Jurisdictions that prefer pure PQC (NCSC UK, CNSA 2.0) accept it as an interim measure. Under OMB M-26-15 the choice is left to the agency's own risk assessment (Appendix D).

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 counterparties support it.

This approach avoids the operational complexity of maintaining different cryptographic configurations per jurisdiction, while remaining compliant everywhere. Track

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

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 the question is no longer hybrid posture but algorithm portfolio: which PQC families the architecture must support, selected per jurisdiction.

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.

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. External specialists augment it for capabilities the organization cannot develop internally in the required timeframe.

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

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
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	Internal procurement with PQC requirements training	0.5

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

Role	Skills Required	Source	Typical FTE
	management; SLA design		
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.

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

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

Skill domain	Primary roles	Proficiency target
		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
Crypto-agility engineering and rehearsal	Crypto-Agility workstream owner (the permanent capability owner after closure), platform engineers	Can execute a rehearsed algorithm change on a Tier-1 system inside the agility window and maintain the mechanisms catalogue as applied to the estate
Cryptographic record ownership	Cryptographic Record owner, CBOM platform team	Can keep the CBOM queryable, fresh under the Activity 2.4 triggers and reconciled to the coverage denominator

Detailed curricula, candidate roles, and evaluation criteria for each domain are maintained in the source article library on PostQuantum.com (Accompanying Resources, front matter).

Build, Borrow, or Buy

Capability	Recommended Model	Rationale

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

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

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

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.
- **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), and delivered as trainings and certifications on these topics by Quantum Academy (QuantumAcademy.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

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

- The Crypto-Agility workstream owner becomes the permanent capability owner, with a standing rehearsal cadence. The Cryptographic Record owner named in Phase 0 keeps the CBOM
- SOC detection rules for cryptographic posture monitoring become part of the permanent detection library

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 three things:

- Whether the system should be running hybrid, classical-only (pending migration), or PQC-only.
- Which specific algorithms and cipher suites are expected.
- 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. GRC governs it. The SOC consumes it. If the registry is kept 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.

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

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 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 key exchange need to know when a flow that should negotiate a hybrid group completes with a classical group instead. An on-path adversary cannot cause that on an intact TLS 1.3 handshake: the transcript that both parties authenticate covers the ClientHello, so a key share removed or altered in transit fails the handshake instead of completing a weaker one. A classical-only session on a hybrid-enabled flow arises in four ways, and detection has to tell them apart:

- **A terminating proxy** or inspection device ends the client's session and opens its own classical connection to the origin.
- **A permissive endpoint policy** still offers or accepts classical-only negotiation.
- **An application retry on a new connection.** A client whose hybrid handshake failed reconnects with a smaller or classical-only offer. Fragmentation-induced failures produce this pattern, and an adversary who can make handshakes fail can force it.
- **An implementation defect** in a library, a product or a configuration.

An adversary can realistically do two of these things: force the third case, or exploit the second. A silent in-transit downgrade of an intact hybrid handshake is not one of the four cases at all. A rule written to catch that will never fire on the mechanism it names.

- **Detection approach.** Write SIEM correlation rules that do two jobs. The first job is to flag the connection: alert on any connection that negotiates a classical-only key share, identified by named group, when the cryptographic posture registry expects hybrid on that flow. The second job is to classify the cause, and there are four causes to tell apart:
  - **A terminating device on the path.** The registry records where each flow's TLS is terminated.
  - **A permissive policy.** The expected configuration still allows classical-only.
  - **A retry.** A failed hybrid handshake from the same client immediately preceded the classical session.
  - **None of these.** Investigate the connection as an implementation defect.

A worked example. Two systems both support ML-KEM, and policy on the flow between them refuses classical-only. A TLS 1.3 session between them should not complete using only ECDH. If one does, there are three possible causes: the policy is not being enforced, a terminating device is on the path, or an implementation is at fault.

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

Each of the three is a finding. Only two of them involve an adversary: a retry that an adversary forced, and a permissive policy that an adversary exploited.

- **Illustrative detection rule.** Alert when both endpoints are listed as "hybrid-enabled" in the cryptographic posture registry, and the TLS handshake between them completes on a key share with no ML-KEM or other PQC key exchange component.

Then correlate that alert with any failed handshake from the same source in the preceding seconds. The correlation separates two cases. A classical session that followed a failed handshake is a retry, and is reported as a retry. A classical session on a first attempt is a policy finding or an implementation finding.

Severity: High.

Measure the false-positive rate locally and use it as a tuning target. This framework publishes no figure for it. Expect the rate to be high wherever the registry is poorly maintained, because stale entries generate constant noise. Note also that an approved compatibility policy that still permits classical negotiation on a flow is not a downgrade, and should not be alerted as one.

- **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 or fail to reassemble the fragmented handshake. The hybrid handshake then fails, and the client retries on a new connection with a classical-only offer, producing a classical session that looks identical to one an adversary forced.

A middlebox cannot strip the hybrid key share from a handshake and let it complete, because the transcript check would fail, so the benign case is always a failed handshake followed by a retry. Detection rules must account for these fragmentation-induced retries, 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 pre-migration backup. A SaaS integration uses an API endpoint that silently reverts to

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

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 negotiating a hybrid or PQC key-exchange group 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.

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

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

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. An adversary who has recovered or emulated a signing key signs outside the HSM. HSM activity monitoring cannot see that forgery. The detection for it sits on the verifier side, and four checks find it:

- an artifact or signature with no matching release record
- signer use outside the release calendar
- a mismatch against the signed release manifest
- an unauthorized trust-store change
- **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: Exfiltration Detection Weighted by Confidentiality Horizon

Passive interception (fiber taps, radio-frequency collection, upstream capture at a carrier or an exchange) produces no signal inside the victim's network and cannot be detected by any SOC control. The control against HNDL is Track A migration, and this use case claims nothing else. The SOC can detect the other route to the same ciphertext, which is exfiltration of stored data from inside the network. That detection is worth

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

more the longer the confidentiality horizon of the data taken, because long-horizon data keeps its value across the decryption timeline.

Corrected at v3.0: earlier editions titled this use case "Enhanced HNDL-Indicator Detection" and reported an "HNDL Exposure Score" built on DLP coverage, which described exfiltration detection as if it observed harvesting. The correction is recorded in the Version History.

- **Detection approach.** The data exfiltration monitoring SOCs already perform, reprioritized by confidentiality horizon. A slow, sustained exfiltration of archived diplomatic correspondence or pharmaceutical R&D data is more damaging than a one-time dump of transactional data, because the former keeps its value for decades.
- **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.** Long-horizon data exposure coverage: for data stores with confidentiality horizons of ten years or more, the share that remain on quantum-vulnerable key establishment in transit (from the CBOM), the share under active DLP and access-anomaly coverage, and the mean time to detect anomalous access to those stores. Target: 100% DLP coverage for stores with 10+ year horizons, mean detection time under 4 hours, and the in-transit share trending to zero on the Track A schedule. The first figure is the migration control and the second is the detection control. Neither is a harvesting indicator.
- **Dependency.** This use case depends entirely on data classification quality. If the organization has not classified its data by confidentiality horizon, the SOC cannot weight alerts by it. 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.

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

- **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 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.

Confirmation that an actor is harvesting a specific organization's traffic is an intelligence outcome obtained from sources outside the organization's own telemetry. No product produces it. A vendor claiming to detect harvesting is detecting something else, because passive collection leaves no trace in the victim's network.

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.

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

- **AI-assisted cryptanalysis as a watch item.** Claims that machine-assisted methods have advanced a classical or quantum attack on a deployed algorithm are tracked alongside hardware progress, under the grading standard below. The operating posture for every strategic source is the same. This framework plans on the assumption that capable actors have at least the public state of the art. That is a planning posture, not a measurement. A published circuit is a conditional estimate for one construction, and the actors with the most capability publish the least.

Indirect indicators are tracked alongside publications: procurement of dilution refrigerators and control electronics, hiring, the cryogenics and control-electronics supply chains, facility construction, and government behavior read as a disclosure in itself. A designated laboratory or government ceasing to publish, or gating disclosure, is the reassessment trigger recorded in Activity 4.7 and never a reassurance.

- **The grading standard for a resource-estimate or break claim.** Grade every claim before it enters the Material Developments KRI. Playbook 1 opens only for a verified claim. There are three grades:
  - **Verified.** The claim has been independently reviewed, either by the authors of the prior estimates or by named reviewers using a disclosed method. Where the claim depends on assumed error rates and code, those have been demonstrated on hardware. The claim states its ledger of assumptions.
  - **Published.** The claim has been posted to arXiv or ePrint. A posting counts as publication, not as review.
  - **Unverified.** The claim's evidence is undisclosed. Any statement of the form "we broke it but cannot show you" grades as unverified, however credible the source.

Two verifications produced with the same model or the same tooling count as one verification. Independence requires that the method and the tooling differ, and that the tooling is disclosed. The tracker records the grade, the reviewers, and the date of each change.

- **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.
- **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.

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

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.

Front-end stage: claims not yet independently verified. This stage applies to cryptanalytic claims against an algorithm or a parameter set. An implementation vulnerability with a vendor advisory or a patch goes straight to the trigger below. Most cryptanalytic claims arrive before verification, and the playbook does not begin with them. Never initiate an emergency algorithm rotation on an unverified cryptanalytic claim. The incident commander named in the playbook may approve proportionate and reversible precautions instead: raising a parameter set where it is already supported, tightening a policy, or accelerating a planned change. Record the basis for the precaution, and the trigger for reviewing it.

The CTI team opens a tracker entry and grades it under the Strategic CTI standard. It then follows three things: the named reviewers, the formal objections, and the open reproduction attempts.

Follow them on a defined cadence. Daily for a claim against a deployed algorithm, weekly otherwise. The 48-hour target is a triage decision, not a completed evaluation. Verification counts as independent only when the method and the tooling differ from the claimant's, and the tooling is disclosed.

GRC issues a communications holding line for board, regulator and counterparty questions. It states what is claimed and what is verified, and the action the organization would take at each grade. The trigger below is a change of grade to verified. A claim that stays unverified is kept under monitoring on the tracker with its review cadence and the reason recorded, because unverified is not disproven.

- **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.

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

- **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.** The SOC detects that a flow the posture registry expects to negotiate hybrid completed a classical-only session instead, and confirms that it is none of the benign cases: not a false positive, not a terminating device on the path, and not a fragmentation-induced retry.

The SOC then attributes the session to one of three causes: a permissive endpoint policy that an adversary exploited, a client that an adversary forced into a retry on a new connection, or an implementation defect (SOC Implementation, Use Case 1).

- **Immediate actions.** Isolate the affected network path. Preserve packet captures for forensic analysis, including any failed hybrid handshake that preceded the classical session. Determine which of the four cases of Use Case 1 applies:
  - A terminating proxy or inspection device that opens its own classical session to the origin.
  - An endpoint whose policy still offers or accepts classical-only negotiation. This includes a compromised endpoint reconfigured to advertise only classical key shares.
  - A client that retried on a new connection with a classical-only offer, after its hybrid handshake was made to fail.
  - An implementation defect.

An on-path device cannot strip the hybrid key share from an intact TLS 1.3 handshake and let it finish. So if both endpoints on a flow refuse classical-only and a classical session completed anyway, the cause is at an endpoint or in an implementation. It is not on the wire. Assess scope: do other flows on the same path, or with the same endpoint, show the same pattern?

- **Escalation.** A confirmed downgrade attributed to an adversary (a forced retry, or an exploited permissive policy) is a strong indicator of an active, sophisticated adversary targeting the organization's cryptographic infrastructure. Escalate to the CISO immediately. A confirmed permissive policy or implementation defect with no adversary behind it is routed to the migration program as an enforcement finding against the agility criterion (Activity 4.6).

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

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.** What the organization migrated beforehand determines the outcome of this playbook. The in-the-moment response changes little. 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.

- **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? Which of the four cases of Use Case 1 was it: a permissive policy the update introduced, a terminating device now on the path, a retry pattern, or an adversary exploiting one of them?

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

- **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 negotiating a hybrid or PQC key-exchange group	Monotonically increasing; regression investigated within 24 hours
Cryptographic Drift Incidents	Systems reverting to classical-only after migration	Zero. Any non-zero value triggers investigation
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
Long-horizon data exposure coverage	In-transit quantum-vulnerable share, DLP coverage and detection latency for 10+ year horizon stores	100% DLP coverage for 10+ year horizon stores; in-transit share trending to zero on the Track A schedule
Algorithm Rotation Drill FP Rate	False positive rate during rotation exercises	Under 10%

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

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.
- **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.

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

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

GRC function produces five instruments. The board uses them to oversee the program, the CISO uses them to manage it, and internal audit uses them to verify it:

- the risk appetite statement
- the KRI framework
- the regulatory intelligence process
- the vendor quantum readiness assessments
- 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.

Its 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 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 the existing one, adapted 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.

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.

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

- **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 and are specific enough to drive decisions, and that they are reviewed annually.

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 at Board, CISO and Operational Levels

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?

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

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
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

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

KRI	What It Reveals	Illustrative Threshold
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)
Unscored Entries	Is the backlog built on classified data?	Count and age of entries flagged "classification missing" (Activity 3.1); any unscored internet- or partner-facing entry, or Tier-1-candidate entry, older than one quarter escalates to the classification owner
Cryptographic Concentration (CCF)	Are our alternatives independent?	Sourced from the Cryptographic Concentration Framework assessment; threshold per the institution's own CCF appetite. No figure computed here

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, permissive policies or misconfiguration	Broken down by the four cases of SOC Use Case 1 (terminating device, permissive policy, retry whether forced or benign, implementation defect) and false positives, with the

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

KRI	What It Catches	Illustrative Threshold
		adversary-attributed subset reported separately
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
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, its deadline archetype and enforcement tier (Regulatory & Standards Alignment Map), 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.

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

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.

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 is required.

Assessment Mechanics

Incorporate quantum readiness into the existing third-party risk assessment process. The assessment should answer four questions about each vendor:

- Has the vendor conducted a cryptographic inventory of the products or services it provides?
- Does it have a PQC migration roadmap with dates?
- Does it support hybrid or PQC cipher suites in current product versions?
- Are its own subcontractors preparing as well?

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.

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

Contractual Provisions

New and renewed contracts with cryptographic dependencies should include four things:

- 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.
- A commitment to support crypto-agile configurations, so that algorithms can change without renegotiating the contract.

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 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, and 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?).

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

How to Evidence

The migration program should produce auditable artifacts at each stage:

- A cryptographic inventory database, with timestamps and coverage metrics.
- CBOM snapshots, quarterly.
- Steering Committee minutes showing KRI review and decision-making.
- The risk appetite statement, with its board approval record.
- The Regulatory Horizon Report, with quarterly updates.
- Vendor assessment records, with remediation tracking.
- The 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 should also observe tabletop exercises (as described in the SOC Implementation section) to assess organizational incident preparedness.

Decision Rights and the Three Lines

The program is first-line management, and it delivers. The risk function is second line, and it challenges the scoring model, the gate decisions and closure. Internal audit is third line, and it provides independent assurance. The board approves the risk appetite and receives the KPI pack. Titles do not move a function between lines: the QRPM and the CISO directing the program are first line for it, and the second-line challenge is owned by a function that does not execute the migration. In a small organization the same people may hold two of the roles in sequence, and the record shows which role they acted in for each decision.

Model Risk in the Phase 3 Scoring Model

The scoring model is a model, and is treated as one. Four disciplines apply to it:

- It is versioned.
- It is documented with its dimensions, its weights, its data sources (the inputs map in Activity 3.1) and its known limitations.
- After each re-scoring run, someone other than the person who ran it validates a sample of entries.
- It changes only by a minuted SteerCo decision that records the weight change and its rationale. The Cryptographic Concentration Framework's independent-review provisions (G.6) are the pattern for institutions that already run a model-risk function.

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

Assurance Questions and the Errors They Prevent

Each question is the test a reviewer applies. The error stated after it is the one the test prevents. The SteerCo asks the same questions before accepting a phase output, so the first, second and third lines apply one test.

- Which findings closed this quarter were classical remediation, and were any counted as PQC progress? Prevents counting classical remediation as PQC progress.
- Which systems reported as migrated show observed negotiation, refusal of classical-only where policy requires it, and a rehearsal record? Prevents reading capability as enforcement, and one migration as the ability to migrate again.
- Which vendor dates in the scorecard are committed rather than forecast, announced or unknown? Prevents accepting a roadmap date as a deployed state.
- What is the coverage denominator, how was it built, and what changed it since the last report? Prevents reading a denominator-driven coverage drop as regression, and endpoint coverage as risk coverage.
- What is the trend per maturity domain across the last two assessments? Prevents reading a maturity snapshot as a trend.
- Which committed dates moved this period, and where is the minuted re-baseline? Prevents treating a governed re-baseline as failure, and an unminuted slip as a re-baseline.
- Which SteerCo decisions this period changed a date, a scope or a risk acceptance? Prevents treating the SteerCo's existence as governance.
- How many entries are unscored, who owns them, and how old is the oldest? Prevents a backlog built on unclassified data.

Insurance Implications

Cyber insurance underwriters are incorporating quantum readiness into their risk assessments. Underwriters are asking about cryptographic inventories, PQC migration plans, and 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."

Qualified Trust Services and Electronic Signature Law

Electronic signature statutes attach legal effect to signatures, and the regimes differ.

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

The EU's eIDAS regulation and Switzerland's ZertES define certificate-based classes, qualified and advanced. A signature's legal weight follows from the assurance around the key, the certificate and the trust service provider.

The US ESIGN Act and UETA are technology-neutral. They attach effect to intent and attribution, without a certificate class. eIDAS already contemplates algorithm weakening: its Article 34 qualified preservation service exists to extend the trustworthiness of a qualified signature beyond the technological validity period, and it is the instrument this framework's re-anchoring strategy maps to.

A qualified signature on a contract, a seal on a filing or a timestamp on a record has legal effect for as long as the document is relied on, which for many records is decades. That reliance period is the verifier's-service-life clock of Activity 1.1, the one Dimension 2 and the Track B timing read. It is neither the record's retention period nor the confidentiality horizon of the data inside it, and a record can be relied on past its retention minimum and retained past the period anyone relies on its signature. Each jurisdiction and service type is assessed separately.

Qualified signatures, seals and timestamps depend on accredited trust service providers. At this edition, this framework has found no provider offering qualified services on post-quantum algorithms, and no published conformity-assessment criteria for them. The search was not exhaustive. Record the position per provider and per jurisdiction, as a dated dependency.

The trust service provider is therefore a Phase 7 vendor-dependency class in its own right. An organization cannot self-issue a qualified signature. Its migration date for statutory signatures is the provider's date.

Ask every trust service provider the four certificate-authority questions of Activity 7.2, add the question of when the provider's qualified status will extend to post-quantum algorithms, and grade the answer for firmness.

TNFL adds a liability question to the statutory one. A signature forged with a future CRQC on a document the organization relied on, or a document the organization signed and a counterparty later repudiates, is litigated on what the organization knew and did at the time. The evidence dossier is the record of diligence, and its weight in a dispute depends on the applicable law and the facts. It shows when the exposure was identified, what the migration plan committed to, which dependencies set the dates (the trust service provider's roadmap among them) and what compensating measures were taken meanwhile. Legal and compliance own this position in the GRC artifact set and review it annually with counsel.

Long-validity signed archives need the re-signing and re-timestamping strategy set out in Appendix H and in Migration Verification & Program Closure. Some archives have a legal effect that must outlive the classical algorithms protecting them. Re-anchor those archives under post-quantum signatures or timestamps before the classical algorithms

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

are deprecated. The archive owner records the schedule in the CBOM, per archive. A qualified timestamp applied now under a classical algorithm is a bridge until a post-quantum qualified service exists, and the record of when and how it was applied is what makes the later re-anchoring defensible.

Crisis Communications and External Stakeholders

A quantum cryptographic event is a confidence crisis before it is a technical incident: operational damage may take days to materialize, but customers, counterparties, regulators, and markets lose trust 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. The evidence dossier should therefore be maintained so that answering is an export and not a project.

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 sector-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.

The GRC function produces four further outputs that the SOC consumes:

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

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.

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

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

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 (BANKING, CAPITAL MARKETS, INSURANCE)

- **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:** In one internal discovery exercise on a limited set of representative components at a point in time, a single mobile banking session invoked ~320 cryptographic function calls, and one payment followed end to end through the interbank chain touched nine parties and 30,000+ unique cryptographic functions. The figures illustrate order of magnitude, not an industry benchmark
- **HSM dependency:** Payment and general-purpose HSM firmware roadmaps constrain migration; payment HSM certification chains, terminal fleets and product concentration are treated in the Payments note below, and HSM model references are in Tool and Vendor References
- **Priority:** Key exchange on external-facing payment APIs; HSM-protected key material; SWIFT interface cryptography

Payments is the specialization of Financial Services for scheme, terminal and payment-HSM constraints. A payments organization uses both notes.

TELECOMMUNICATIONS

- **Scale:** 5G networks involve cryptography at every layer from radio (5G-AKA) to core (service mesh, IMS, lawful intercept)

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

- **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 (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
- **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: January 1, 2027 for new NSS acquisitions, December 31, 2030 for phasing out equipment that cannot support CNSA 2.0, December 31, 2031 for mandated use of CNSA 2.0 algorithms (CNSSP 15, updated 2024 and published March 2025, with stated exceptions). The advisory's per-category chart sets its own dates and does not extend that schedule: 2030 for software/firmware signing and networking equipment (VPNs/routers), and 2033 for web/cloud/OS platforms, for niche equipment, and for updating or replacing custom applications and legacy equipment, which is the chart's last date and falls after the December 31, 2031 mandated-use date.

A system's binding date is the CNSSP 15 date or the date its applicable profile notes. ML-DSA-87 is the CNSA 2.0 signature algorithm for any use case, including firmware and software signing; LMS and XMSS are allowed for signing firmware and software, with LMS SHA-256/192 recommended and the transition preferred to begin with them,

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

while HSS and XMSS^MT are not allowed for NSS; SLH-DSA is not approved for any NSS use, and HashML-DSA is not allowed at present (NSA CNSA 2.0 FAQ, Version 2.1, December 2024).

- **Classification constraints:** Working with classified systems adds complexity for testing and deployment
- **FedRAMP/FIPS requirements:** In US federal and defense systems the validation question is asked per operation, not per product. Signing and key generation, verification, and key establishment are recorded as separate operations, each read against the authority that binds the system: FIPS 140-3 validation where an agency policy, a contract or a protection profile requires a validated module, and CNSSP 15 with the applicable CNSA 2.0 profile for national security systems.

Key generation and signing must enter production through a module whose certificate lists the operation as an approved service, and SP 800-208's conditions for stateful hash-based signing keys apply unchanged. NSA's CNSA 2.0 FAQ treats software and firmware signature verification differently: for a verification-only implementation, CAVP algorithm validation is sufficient in the case the FAQ states. Validated modules with post-quantum approved services exist in software (certificate #5247, ML-KEM at ML-KEM-768 and ML-KEM-1024 only, April 2026) and in Level 3 hardware (certificate #5497, ML-KEM, ML-DSA, SLH-DSA and LMS, with HSS vendor-affirmed, August 2026, with the Security Policy qualifications stated in Phase 5); coverage for any other product is read from its own certificate, and the Deployment Environment Classification in Phase 5 sets the per-product rule. Module and firmware references are in Tool and Vendor References.

- **Priority:** Comply with CNSA 2.0 milestones; ensure PQC-capable products in procurement pipeline

PAYMENTS

- **Highest HNDL urgency:** Cardholder data and payment credentials cross many parties per transaction, and account numbers and track data keep their value to an attacker for the life of the card and beyond (a statement about the attacker's horizon, not a license to retain them: PCI DSS prohibits storing sensitive authentication data after authorization)
- **Scheme-mandated timelines:** Card-scheme and payment-system-operator requirements arrive as market-access gates (Enforcement Tiers) and set dates that merchants, acquirers and processors cannot move
- **HSM certification chains:** Payment HSMs carry PCI PTS HSM certification as well as FIPS 140-3, and a PQC firmware that breaks the PCI certification is not deployable until re-certified. The Activity 6.2 register records both certificate references

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

- **Terminal fleets:** POS terminals and PIN pads have multi-year lifecycles, in-field key-injection and remote-update constraints, and vendor-controlled cryptography; treat them as the long tail of the payments estate, the way OT is treated above
- **HSM vendor concentration:** A small number of payment HSM product lines sit under most of the sector; whether two acquirers' alternatives share one firmware family is a Cryptographic Concentration Framework question
- **Priority:** Key exchange on external-facing payment APIs and scheme interfaces; HSM firmware and certification roadmaps; terminal-fleet refresh alignment (Payments Extension)

DIGITAL ASSETS

- **On-chain public-key exposure:** A public key that has become visible on a ledger, from the address format or once the address has signed a transaction, is harvestable for future forgery, and for most account models there is no re-key without moving the asset; inventory addresses by exposure state
- **Immutability of forged transactions:** A transaction signed with a forged key is final once confirmed, so TNFL is the primary threat and recovery is a governance question rather than a cryptographic one
- **Protocol-governance dependency:** The signature algorithm is set by the protocol, so migration on the asset side waits on protocol-level change the organization does not control; treat the protocol like a counterparty under Activity 7.6
- **Custody and signing infrastructure:** Custodial key management, HSM-backed and multi-party signing stacks are the migratable surface today. The shared-verifier-defect failure mode is treated in the Digital Assets Extension
- **Priority:** Custody key management and signing stacks; address-rotation policy for exposed keys; protocol roadmap tracking (Digital Assets Extension)

© 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. The phase text points to any related appendix, and all of them are also usable as standalone reference cards for practitioners who are already familiar with the framework. Appendix I contains invented, labeled worked artifacts across the phases; Appendix J is the requirements register, generated from the text at each build.

© 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	Hash-based signature	Digital signatures (conservative	Varies widely by parameter	Stateless hash-based; no lattice assumptions.

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

NIST Standard	Algorithm	Type	Primary Use	Key/Signature Size	Notes
	SPHINCS+)		choice)	set	Larger signatures but different security foundation. Recommended as backup/alternative if lattice cryptanalysis advances.

Pending standardization (NIST):

- **FN-DSA (formerly FALCON):** FIPS 206 is under development, with no published draft and no published date: NIST stated on April 2, 2026 that the draft was still in clearance within NIST and the Department of Commerce, and as of September 13, 2026 FIPS 206 appears in NIST's FIPS listing in neither final nor public-draft status, so secondary reports of a published initial public draft are wrong. Compact lattice-based signatures suitable for bandwidth-constrained environments, though signing requires floating-point arithmetic that must be implemented in constant time. Expected to complement ML-DSA where signature size is critical; not the HSM default (Two-Track Migration Model, Phase 5).
- **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 FIPS 207 is under development, with no published draft and no published date.

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)

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

Standards Pipeline (September 2026)

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. Under development at NIST at this edition (September 2026), with no published draft and no published date (Appendix A above states the evidence). 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 mathematical hardness assumption (error-correcting codes versus structured lattices). If a future attack compromises lattice-based KEMs, HQC provides a non-lattice fallback. Selected for standardization; FIPS 207 is under development, with no published draft and no published date, and as of September 13, 2026 it appears in NIST's FIPS listing in neither final nor public-draft status.

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 14, 2026 that nine candidates had advanced: FAEST, HAWK, MAYO, MQOM, QR-UOV, SDitH, SNOVA, SQIsign, and UOV. The HAWK team withdrew its algorithm on July 29, 2026, leaving eight in evaluation. Evaluation is expected to take approximately two years, so these candidates are not for production planning. Their one use is 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 for any use case, including signing firmware and software, use ML-DSA-87 (NIST Security Level 5). LMS and XMSS (SP 800-208) are allowed for the specific application of signing firmware and software, with all parameters approved and LMS SHA-256/192 recommended; HSS and XMSS^MT are not allowed for NSS.

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

SLH-DSA is not approved for any NSS use, and HashML-DSA is not allowed at present (NSA CNSA 2.0 FAQ, Version 2.1, December 2024). Symmetric cryptography remains AES-256 and hash functions remain SHA-384 or SHA-512, unchanged from CNSA 1.0.

Some financial services organizations adopt the CNSA 2.0 parameter sets (ML-KEM-1024, ML-DSA-87) as a single internal standard even when not required, because one parameter set simplifies procurement across government and commercial contexts. The choice is an internal policy, not a compliance conclusion.

© 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. The first-layer table in Activity 5.2 picks the layer per architecture pattern.
- **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 (Program Foundations). 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 Inequality 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 need, read from the confidentiality horizon of Activity 1.1 and never from its retention period. For TNFL-exposed signatures, X is the period over which the signature will still be relied on: the verifier's service life for firmware and fielded devices, the reliance period of a signed record, or the anchor's remaining period of authority where every relying party can be updated when it ends, whichever is longest for the anchor. These are the clocks Dimension 2 reads.

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. Even the low ends of those ranges put X + Y at 14 years against a Z that may be 10, and each year of delay reduces Z by a year without reducing X or Y.

© 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), as a recommendation. The certification "visa" is ANSSI's instrument (next row), not BSI's	Hybrid should be the default design assumption in German-facing deployments
ANSSI (France)	**Strongly recommended** in the short and medium term (ANSSI's views on the post-quantum transition and its follow-up position paper, framed as recommendations and not as a general obligation on French organizations); hybridation is mandatory for products seeking ANSSI's security visa (*visa de sécurité*) at the phase-2 evaluation	Hybrid is the safer default for French assurance-sensitive products, and a requirement for a product seeking the security visa
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	Flexibility to choose hybrid or pure PQC depending on the

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

Jurisdiction	Hybrid Position	Implication for Practitioners
	800-227 discusses X-Wing as an example of a general-purpose hybrid KEM	system
OMB M-26-15 (US federal agencies, June 24, 2026)	**Agency risk decision**: hybrid architecture described as a useful tool for managing risk during migration and as an "intricate and resource-intensive stopgap" that introduces its own risks and complexities; agencies evaluate its tradeoffs thoroughly and may implement it on their own risk assessments and technical requirements	Hybrid is permitted, evaluated and recorded per deployment under the operative US civilian guidance. It is not a default the guidance sets
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

Registry markings. In the IANA TLS Supported Groups registry, X25519MLKEM768 is Recommended=Y (RFC 10024, standards track); SecP256r1MLKEM768 and SecP384r1MLKEM1024 are Recommended=N and are the NIST-curve combinations (in all three, the approval comes from the ML-KEM component in a validated module; the NIST-curve groups add an approved classical component); Recommended=N is not a prohibition; the pure ML-KEM groups are Recommended=N in an Informational document in the RFC Editor queue at this edition; and the pre-standard Kyber768 draft groups are obsoleted at Recommended=D.

A jurisdiction that prescribes pure ML-KEM is asking for a Recommended=N group, which is permitted and is recorded as a compliance profile. The registry marking is what an auditor reads.

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

Multinational recommendation: Deploy hybrid as the default deployment pattern. This satisfies the strictest positions in the table (the continental European recommendations and the ANSSI security-visa requirement), 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 is accepted in every jurisdiction in the table when confirmed per deployment profile against the applicable authority; under OMB M-26-15 a US federal agency adopts it on its own risk assessment and records the tradeoff evaluation, rather than as a default the guidance sets.

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

APPENDIX E: QUICK-REFERENCE CHECKLISTS

90-Day Quick Start Checklist

- ☐ Executive sponsor identified and committed
- ☐ QRPM appointed
- ☐ SteerCo convened; charter published; RACI defined
- ☐ Multi-year budget (minimum 3 years) approved or in approval process
- ☐ Top 20 critical systems identified (Phase 0 scoping)
- ☐ Cryptographic discovery tool selected and procurement initiated
- ☐ Quantum risk added to enterprise risk register
- ☐ Cryptographic policy updated to include PQC-approved cipher suites
- ☐ Procurement policy updated to require PQC roadmaps from vendors
- ☐ Training cohort (10–20 people) enrolled in PQC fundamentals
- ☐ First 2 hybrid pilot targets selected
- ☐ First rehearsal target selected (Activity 5.2)
- ☐ Vendor questionnaires sent to top 10 strategic vendors
- ☐ KPI baseline values established; Q+1 targets set
- ☐ 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 verified in use with post-quantum signatures at their relying parties (target: Y); readiness line: X anchors PQC-capable and not yet verified in use; lifecycle line: X anchors shortened or rotated this quarter
- **Readiness evidence, agility KPIs (Metrics, KPIs & Reporting):** configuration-driven share of Tier-1 systems [X%, target Y%]; rehearsed time-to-change on the most recent Tier-1 rehearsal per track [X hours against the agility window]; agility-debt burn-down [N entries closed this quarter; N remaining]; new-system agility compliance [X%]
- **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 companion book, *Quantum Ready* (QuantumReady.com).

"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 migration duration and Z time until CRQC) shows that organizations with long-lived sensitive data and multi-year migration timelines must start now regardless of the CRQC date. HNDL is treated as an active present-tense risk on the basis stated in the Two-Track Migration Model: the means, motive, opportunity and precedent for passive collection are demonstrated.

The harvest itself is unobservable by design. For the same reason, the objection that no harvesting has been shown fails: a passive collector leaves no trace in the victim's network, so finding none is expected and is not reassurance. The July 2026 PostQuantum.com analysis of why HNDL cannot be detected sets the reasoning out in full for use in the funding conversation.

Three published statements are consistent with treating HNDL as a present risk: CISA, NSA and NIST wrote in their joint factsheet of August 21, 2023 that threat actors "could be targeting data today that would still require protection in the future". The Dutch AIVD, CWI and TNO wrote, publishing the second edition of their PQC Migration Handbook on December 3, 2024 (first edition April 2023), that secured data can be "intercepted today and then deciphered with a quantum computer from Q-Day onwards"; and Mascelli and Rodden describe HNDL in Federal Reserve Board staff working paper FEDS 2025-093 (September 2025) as "a present, active, and in some circumstances unavoidable data privacy risk," under the paper's own disclaimer that the views are the authors' and not the Board's.

Every day without hybrid key exchange creates permanent exposure. Binding dates now exist, and the earliest of them fall at the end of this decade.

NIST IR 8547 proposes deprecating 112-bit-strength quantum-vulnerable public-key algorithms after 2030, and disallowing all quantum-vulnerable public-key algorithms after 2035. That is the reference point most other instruments are built against. The

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

roster by jurisdiction is in the Regulatory Timeline Context, and its current version is on the Global PQC Migration Clock page.

"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.

"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 inside broader IT modernization that is already justified on its own terms. Four examples:

- PKI automation, for shorter certificate lifetimes.
- FIPS 140-3 migration, for the FIPS 140-2 sunset.
- HSM refresh, for end-of-life hardware.
- 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

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

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." Supervisors now publish readiness indices and sector surveys, so this claim is checkable against data the regulator holds, and the data does not support it. The Hong Kong Monetary Authority's Quantum Preparedness Index (July 2026) stages banks across awareness, planning, pilots and practical preparedness. Its first reading reported roughly half of surveyed banks without a formal post-quantum plan, and roughly a third exploring or piloting quantum-related initiatives of any kind.

FINMA's July 2026 survey of Swiss institutions reported 8% with a specific roadmap, against 72% that had neither planned nor implemented quantum-safe measures.

On the deployment side, a June 2026 internet-wide measurement by Forescout Research – Vedere Labs, used here as a data source, found PQC-capable SSH on 11.8% of servers, and inside enterprise networks on 50% of IT devices, 28% of IoT devices, 16% of OT devices and 6% of medical devices (device and version capability in the measured population, not the share of institutions that have migrated). Each figure is the answer to a different question: the supervisory surveys measure institutions' planning, the device measurement measures capability in an observed population, and none of them is a migration rate or a sector median.

Client-side PQC is enabled by default across the major browsers and runtimes (PQC Deployment Status, Phase 5). An organization that is waiting for the sector is answering to a supervisor that has already asked which stage of the index it has reached.

"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.

"We will wait for FIPS-validated modules." The validation gap is real, and this framework treats it as a hard constraint (Deployment Environment Classification, Phase 5). Validated modules with post-quantum operations as approved services exist in software (certificate #5247, April 2026, ML-KEM at ML-KEM-768 and ML-KEM-1024 only) and in Level 3 hardware (certificate #5497, August 2026, ML-KEM, ML-DSA, SLH-DSA and LMS, with the qualifications stated in Phase 5), and FIPS-required environments run PQC in approved mode only through the module whose certificate covers the operation.

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

The constraint decides which operations run in approved mode, in which environments. It does not decide whether the program starts.

Seven activities run without a validated module at all: discovery, the CBOM, risk scoring, the roadmap, vendor engagement, certificate-lifetime reduction and the HSM register. Unrestricted and FIPS-aware environments deploy hybrid key exchange now. FIPS-required environments plan, stage and rehearse, so that the day a certificate issues is a change window rather than a project start. A program that waits for the certificate arrives at it without any of the work that makes the certificate usable.

"We bought a quantum-safe product, so we are covered." Decompose the claim into what is covered and what is not, per the claim-evidence principle. Which functions: key exchange, signatures, or both? Which layer: the product's own TLS, or the keys, certificates, firmware signing and identity stack around it? Validated how: a configuration option, a CAVP-tested algorithm, or a CMVP-validated module, and for which version and operating environment?

What remains: every flow that does not pass through the product, every counterparty on the far side of it, and every signature it verifies against a classical trust store. A vendor that has updated its TLS library while HSM firmware, certificate management and key-lifecycle tooling stay classical has done a fraction of what the migration requires, and the CBOM entry records how much. The purchase is recorded as one migrated component under the evidence standard of Migration Verification & Program Closure, once the negotiated outcome is observed. It is never recorded as coverage.

© 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
Phase 7 – Vendor & Supply Chain	Govern (GV.SC)	Planning and Execution	Stages 2–3 (supplier actions)	Planning / Execution

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

Framework phase	NIST CSF 2.0	PQCC Roadmap	ETSI TR 103 619	Dutch Handbook
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:

CNSA 2.0 milestone	What meets it
The 2027 acquisition requirement	Phase 7 procurement language
The 2030 software and firmware signing deadline	Track B's SP 800-208 deployment
The 2030 networking equipment deadline	Track A's infrastructure waves
The CNSSP 15 phase-out of December 31, 2030 and use mandate of December 31, 2031	The roadmap's Year 4 to 5 waves for NSS
The advisory chart's 2033 horizon	The outer years of the multi-year roadmap

Supervisory Readiness Indices

Supervisors have begun publishing readiness indices and sector surveys that stage institutions along an adoption path. The crosswalk maps each index's stages to the phases that produce the evidence for them. Scores and dates belong on the Global PQC Migration Clock page, not here. Successor indices are added as they publish.

Index	Stage or dimension	Framework phases that produce the evidence
HKMA Quantum Preparedness Index (banking sector, July 2026)	Awareness	Phase 0 (board briefing, mandate, risk-register entry)

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

Index	Stage or dimension	Framework phases that produce the evidence
HKMA Quantum Preparedness Index	Planning	Phases 1 to 4 (inventory, CBOM, scoring, roadmap and governance)
HKMA Quantum Preparedness Index	Pilots	Phase 5 (pilots and rehearsals) and Phase 6 (infrastructure tests)
HKMA Quantum Preparedness Index	Practical preparedness	Phases 5 to 7 at scale; Migration Verification & Program Closure; the Crypto-Agility foundation
Supervisory sector surveys (FINMA and successors)	Planned; roadmap in place; measures implemented	Phase 4 roadmap; Phase 5 verification records; the evidence dossier

Sibling Instruments

The family's sector taxonomy is this framework's, and the Cryptographic Concentration Framework's Payments and Financial Services extensions align to it.

Instrument	Its structure	This framework's phases
Applied Quantum CBOM Profile (v1.0-RC or later)	Tiers 1 to 3: the minimum record, the evidence and roadmap fields, the custody fields	Phases 1, 2, 3 and 7, as the informative machine-readable carriage of the records those phases define
Applied Quantum CBOM Profile	Tier 4	Consumed by the Cryptographic Concentration Framework, not by this framework
Cryptographic Concentration Framework (CCF)	Its own Phases 0 to 7, a second-line instrument for risk analysts	Fed by this framework's Phases 1, 2 and 7 (inventory, CBOM, vendor register); consumed at Phase 3 as scoring context, at Phase 7 in alternative-sourcing

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

Instrument	Its structure	This framework's phases
		decisions, and in the CISO-level KRI cascade

Crypto-Agility Maturity Models

Model	Levels	This framework's Crypto-Agility domain
Crypto-Agility Maturity Model (CAMM; Hohm, Heinemann and Wiesmaier, Hochschule Darmstadt, 2022)	Five levels, 0 to 4, each defined by a set of requirements drawn from the crypto-agility literature. A system reaches a level when it meets that level's requirements	Levels 1 to 5 of the Maturity Model, with the rehearsal criteria of the Crypto-Agility foundation as the measurement method

© 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 uses quantum-vulnerable cryptography, and several are routinely forgotten because no single team owns them. Deployability reflects the state of standards and implementations as checked for this edition (June 2026 data, re-verified in September 2026 where a standard moved) 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)	RFC 10042 (August 2026, Informational) for PQ/T hybrid key exchange with ML-KEM; 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, OpenPGP)	TNFL	S/MIME: RFC 9881 and RFC 9882 (ML-DSA in X.509 and CMS, October 2025), RFC 9814 (SLH-DSA in CMS). OpenPGP: RFC	Early adoption with caveats: S/MIME as issuers and clients support the profiles; OpenPGP once interoperability	Track B; Activity 5.6 for archives

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

Protocol family	Quantum threat	Standards status (June 2026)	Deployability	Framework reference
		9980 (June 2026, Proposed Standard), composite ML-KEM and ML-DSA with elliptic-curve components, SLH-DSA standalone, no RSA composite	is tested across the two implementation dialects in the field, and the installed RSA key base re-keyed because the RFC gives it no upgrade path	
DKIM	TNFL (RSA and Ed25519 signing keys, RFC 8463)	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	No unilateral algorithm change possible; inventory signers and validators, build algorithm agility into the signing pipeline, test rollovers in private zones	Track B
Code and firmware signing	TNFL (highest persistence)	SP 800-208 LMS/XMSS final; composite ML-DSA drafts	Deploy now where the signer meets SP 800-208 section 8 (Level 3 hardware, no key export, in-module state), the module's certificate names the parameter set in use, state safety is accepted on	Track B; Activity 5.2

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

Protocol family	Quantum threat	Standards status (June 2026)	Deployability	Framework reference
			evidence inside the implementation (supported parameter sets, exhaustion monitoring, atomic state handling, tested backup, replication, crash-recovery and concurrent-signing behavior) and the verifiers accept the parameter set; dual-signed under a verifier policy that requires both	
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; qualified trust services cannot yet be satisfied with PQC (GRC Implementation, Qualified Trust Services and Electronic Signature Law)	Track B; Closure
Long-lived signed records	TNFL	Re-sign versus Merkle-anchor strategies;	Decide per archive; document in the	Activity 5.6; Closure

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

Protocol family	Quantum threat	Standards status (June 2026)	Deployability	Framework reference
and archives		evidence record syntax	CBOM	
Kerberos / Active Directory (PKINIT)	TNFL; HNDL on ticket exchange	Vendor-dependent; early	Track operating-system vendor roadmaps	Track B identity
FIDO2 / WebAuthn passkeys	TNFL (typically ES256 or RS256 credentials; inventory the enrolled algorithms)	PQC COSE codepoints assigned; ML-DSA for COSE is RFC 9964 (May 2026); WebAuthn integration and authenticator support in progress	Monitor; plan re-enrollment; ask CIAM vendors	Track B identity
OIDC / SAML token signing	TNFL (JOSE RS256/ES256 for JWT; XML Signature RSA/ECDSA for SAML)	Algorithm-agile by design; ML-DSA for JOSE is RFC 9964 (May 2026); identity-provider support 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 and personal)	HNDL (SAE key establishment in personal mode); TNFL (enterprise certificates via EAP-TLS)	Enterprise mode rests on EAP-TLS or another EAP method and follows TLS; SAE (personal mode) is a password-authenticated key exchange over elliptic-curve or finite-	Cover enterprise via the 802.1X row; treat SAE networks carrying sensitive traffic as HNDL-exposed until a standard exists	Phase 6

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

Protocol family	Quantum threat	Standards status (June 2026)	Deployability	Framework reference
		field groups, quantum-vulnerable at key establishment, with no post-quantum SAE standardized		
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.

APPENDIX I: WORKED EXAMPLES ACROSS THE PHASES

Every artifact below is invented. The organization, the names, the dates and the figures are illustrative and refer to no real program. The appendix is informative under the Normative Convention. Each artifact is what "done" looks like on paper for one phase or for closure, in the form the relevant Outputs table asks for. Phase 3 has two (a scored entry and an appeals record) and Phase 4 has two (a gate record and a re-baseline record), because they are different decisions.

Artifact	**Phase**	**The Outputs-table row it instantiates**
I.1 Charter purpose statement	Phase 0	Approved program charter
I.2 Priority A scoping decision	Phase 1	Risk-driven scoping document
I.3 Scored CBOM entry	Phases 2 and 3	CycloneDX CBOM; scored and tiered CBOM
I.9 Appeals record	Phase 3	Scored and tiered CBOM (appeals under the ratified rule)
I.4 Gate record	Phase 4	Gate records (Activity 4.6)
I.8 Re-baseline trigger record	Phase 4	Re-Baseline Protocol record (Activity 4.8)
I.5 Wave rehearsal record	Phase 5	Rehearsal record (Activity 5.2)
I.6 HSM register row	Phase 6	HSM/KMS upgrade schedule (the register, Activity 6.2)
I.7 Questionnaire response graded for firmness	Phase 7	Vendor questionnaire responses; vendor PQC readiness register
I.10 Closure decision record	Migration Verification & Program Closure	Closure decision record

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

I.1 Phase 0 – Charter Purpose Statement in the Dual-Outcome Form

- **Program:** Quantum-Safety Readiness Program, Halden Financial (invented). **Sponsor:** Chief Information Security Officer. **QRPM:** appointed March 2, 2026.
- **Purpose.** The program delivers quantum-safety readiness against the organization's dated obligations (NIST IR 8547's proposed 2030 deprecation of 112-bit-strength public-key algorithms and 2035 disallowance of all quantum-vulnerable public-key algorithms as the reference deadlines; the payment scheme's March 2028 procure-by expectation for external-facing APIs at enforcement tier 1; the EU coordinated roadmap's high-risk migration milestone as the supervisory expectation), and it delivers that readiness through crypto-agility: the PQC migration is the first exercise of a standing capability to change the organization's cryptography on a planned date, and the capability is what the program leaves behind.
- **Success criteria.** (1) No new HNDL exposure on Tier-1 data in transit by December 2027: hybrid key exchange observed in production on every Tier-1 internet-facing and partner-facing endpoint, classical-only refused where policy requires it. (2) Long-lived trust anchors protected: root and intermediate lifetimes shortened by December 2026; SP 800-208 firmware signing deployed on the two in-house signed platforms by June 2027.

(3) Regulatory deadlines met with the twelve-month buffer the risk appetite statement sets. (4) Crypto-agility embedded as organizational capability, proven at closure by a rehearsed second algorithm change on at least one Tier-1 system per track inside the agility window, with the agility-debt register at zero unaccepted debt.

- **What the program is not.** It is not a project that ends when the last Tier-1 system is migrated. The Crypto-Agility workstream and the Cryptographic Record transfer to named permanent owners with funded budgets at closure, and the charter names both owners today.

I.2 Phase 1 – Priority A Scoping Decision

- **Decision record, Activity 1.0.** March 20, 2026. Prepared by the Discovery workstream lead; approved by the QRPM.
- **Rule applied.** Priority A = internet-exposed or partner-exposed, and either handles data with a confidentiality horizon past 2035 or carries a trust anchor with validity past 2030.
- **Inputs.** Asset register export (4,212 systems); cloud account list (three providers, 61 accounts); data classification register (the Sensitive Data Inventory of Activity 1.1). Two coverage denominators, kept separate: 4,212 systems from the register, and 61 cloud accounts for account enumeration; systems discovered inside those accounts are reconciled to the register and given a stable identity before the system denominator changes, and the change is recorded with the next coverage report. Neither denominator comes from a tool's count.

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

- **Result.** 118 systems in Priority A: 47 internet-facing customer and partner endpoints (the customer web and mobile API gateway, the open-banking API gateway, the file exchange with three processors); 29 partner-facing interconnects; 12 internal systems holding long-horizon data reachable from the internet-facing tier; 30 trust-anchor systems (two internal CA hierarchies, the code-signing service, the HSM cluster, the identity provider's token-signing keys). Priority B: 640 systems. Everything else is Priority C, found by continuous discovery on its own schedule.
- **Accepted for now.** Three partner interconnects whose traffic cannot be tapped at a choke point the organization controls are on the accepted-gaps register, owner the network lead, closure date June 30, 2026.
- **Why the record exists.** The first question an auditor asks of Phase 1 is why these 118 and not others. The rule, the inputs and the count answer it.

I.3 Phases 2 and 3 – One Scored CBOM Entry

Minimum record (Activity 2.1). Component identifier: api-gw-prod-01 (CMDB CI-20831). Algorithm: ECDHE on secp256r1 (P-256) for TLS 1.3 key establishment; curve OID 1.2.840.10045.3.1.7 (the curve identifier, recorded separately from the TLS group and from the certificate signature algorithm). Key size / security parameter: 256-bit curve; quantum-vulnerable. Protocol context: TLS 1.3, server side, open-banking API gateway, partner-facing.

Implementation: OpenSSL 3.2.1 in gateway appliance firmware 7.4. Certificate reference: cert-4471, ECDSA P-256 leaf, expires November 2, 2026 (scored as its own Track B entry). Cryptographic function: key establishment (Track A). Confidentiality horizon: 20 years (account and transaction data). Regulatory deadline: December 31, 2035, the NIST IR 8547 (draft) disallowance date for a 128-bit-strength curve, enforcement tier 4, as the reference (the draft's 2030 deprecation applies to 112-bit strength and is not this entry's milestone). The scheme's March 2028 procure-by expectation, tier 1, recorded from the scheme's technical requirement.

Data classification: confidential, regulated. Quantum vulnerability status: vulnerable (Shor), ECC attack-cost band. Migration status: planned. Owner: Payments Platform team. Vendor dependency flag: vendor-dependent (appliance firmware); roadmap firmness: forecast (firmware 8.0, Q1 2027, stated in writing, no remedies). Evidence grade: E1, runtime observed by passive monitoring, March 14, 2026.

Phase 3 scoring (ratified model v1.0; Low = 1 to Critical = 4). Dimension 1, HNDL: horizon over 15 years, Critical; accessibility partner-facing, High; volume (high-volume API traffic) recorded as the sequencing tie-break and not scored: 4. Dimension 2, TNFL: key-establishment entry, not applicable, excluded. Dimension 4, Regulatory: a tier-1 instrument within two years, Critical: 4. Dimension 3, Feasibility: vendor-dependent with a forecast date, Hard; inverse score 3.

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

Tier, by the ordered test of Activity 3.2: Tier 1 on the first condition, Dimension 1 Critical with accessibility High (a partner-facing entry with a 20-year horizon). The score plays no part in that assignment. Sequencing score: Priority = ((4 × 0.35) + (4 × 0.25) + (3 × 0.15)) / (0.35 + 0.25 + 0.15) = 2.85 / 0.75 = 3.80, the maximum reachable over these three applicable dimensions (the four-dimension maximum is 3.85). The score is the entry's position among the partner-facing Tier-1 entries in the Track A sequence, ahead of any partner-facing Tier-1 entry scoring below 3.80.

The migration is made possible on the program's timeline by the bridging pattern of Activity 7.4 (a PQC-capable proxy in front of the appliance until firmware 8.0 ships). Recorded at Phase 3, the five Activity 2.1 fields marked for this phase: trust scope and verifier updatability not applicable to a key-establishment entry; dependency firmness forecast (firmware 8.0, Q1 2027); delivery state blocked on the appliance firmware, with the confidentiality leg made available by the bridging pattern, so the entry stays Tier 1 and appears on the roadmap against the firmware date; compensating control the bridging proxy, with the classical path behind it recorded as exposed until the firmware ships. Agility-debt register: listed, because the algorithm changes only with a firmware release.

I.4 Phase 4 – Gate Record with the Agility Criterion

- **Gate G5 → G6, SteerCo, November 12, 2026.** Decision authority: SteerCo. In scope: three pilots (customer web edge; site-to-site VPN; internal service mesh).
- **Criteria.** Pilots validated: yes, three pilot reports filed. Performance baselines established: yes; p99 handshake latency at the edge +11 ms, inside the SLO. Agility criterion on every pilot: negotiated outcome observed in production traffic, yes on all three (passive-monitoring records October 28 to November 9); classical-only refused where policy requires it, yes on the mesh, not required at the edge during the published dual-stack window (recorded); algorithm changeable by configuration, yes on the edge and the mesh, no on the VPN concentrator (parameter set fixed until firmware 9.2, listed as agility debt); change rehearsed with rollback, yes for the edge and the mesh (rehearsal records R-01 of October 22 and R-02 of October 30, with the R-02 retest of November 5), not rehearsed on the VPN.
- **Protection result.** Passed on all three pilots: hybrid negotiation observed in production on each, classical-only refused where policy requires it. **Agility result.** Passed on the edge and the mesh (rehearsed, rolled back, configuration-driven); agility debt on the VPN concentrator, closure date the firmware 9.2 release. **Decision.** G6 passed for the edge and the mesh; Tier-1 rollout approved for both patterns. The VPN pilot exits as migrated with agility debt and does not enter Wave 2 until the rehearsal is completed on firmware 9.2 or a compensating control is accepted by the risk owner. Minuted; filed to the evidence dossier as GATE-05.

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

I.5 Phase 5 – Wave Rehearsal Record

- **Rehearsal R-02, Wave 2, internal service mesh, October 30, 2026; retest November 5, 2026.** Change rehearsed: key-establishment group X25519MLKEM768 to SecP384r1MLKEM1024 on the mesh's mTLS policy, in staging, rolled back at the end.
- **Elapsed times.** Assessment (trigger to the decision that a change is needed, using the 48-hour assessment test): 6 h 40 min. Decision (approval under the standing change authority): 19 h 15 min, waiting for the next change-advisory window. Change (policy rollout to 412 staging workloads): 1 h 05 min. Verification (new group observed in staging traffic; rollback exercised and observed): 2 h 30 min. Total for the first attempt: 29 h 30 min, inside the agility window (the 48-hour assessment plus the organization's ten-business-day change window); time to full acceptance, including remediation of the two workloads and the retest: six calendar days, also inside the window.

The finding the record exists to show is not a breach of the window: it is the decision segment, which took two thirds of the first attempt, and the two workloads that pinned the group, which are agility debt found in staging before production found it.

- **Findings.** The decision segment dominates: cryptographic policy changes move from the weekly change-advisory board to a pre-approved change class, which is the next investment. 410 of 412 workloads followed the policy. Two pinned the group in their own configuration, were assigned remediation owners, and passed a retest of change and rollback on November 5, so the record shows 412 of 412 only after the retest. Rollback restored X25519MLKEM768 on all workloads in 40 minutes.
- **Signed:** Crypto-Agility workstream owner; mesh platform lead. **Feeds:** rehearsed time-to-change KPI, December report.

I.6 Phase 6 – HSM Register Row

Field	Value
Visible platform	Network HSM appliance, model as procured, serial 7A-2210, data-center A partition set
Supplier's canonical firmware family	Family L7, release 7.9
Manufacturer of origin	The supplier itself; not a rebadged OEM unit
FIPS 140-3 certificate reference	EXAMPLE-CMVP (a placeholder; no real validation is asserted), module version 7.9.0, tested operating environments as listed on the certificate; ML-KEM and ML-DSA present as non-approved services at this firmware;

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

Field	Value
	approved-mode PQC not validated
Entropy	ESV certificate EXAMPLE-ESV (a placeholder), physical noise source, operating envelope 5 to 40 °C as stated on the certificate; evaluated at multiple operating conditions: no, the validation predates IG D.K Resolution 23, and the next procurement asks for it
PCI HSM certificate	Not applicable, general-purpose HSM
Standing failover exercise	Quarterly failover to the backup partition, last executed September 15, 2026, 24 minutes; recorded as a failover and state-management result, not as an algorithm-change rehearsal, and its timing does not feed the rehearsed time-to-change KPI (Activity 5.2)
Dependency boundary	On-premises. The cloud KMS rows of this register are boundary rows, not manufacturer rows

I.7 Phase 7 – Questionnaire Response Graded for Firmness

- **Vendor:** the API gateway appliance supplier, Strategic Blocking. Questionnaire returned May 8, 2026.
- **Configuration-driven algorithm change:** "Yes from firmware 8.0; in the 7.x line the key-exchange groups are fixed per release." Graded: capability announced. The current release fails the criterion. **Updatable verifiers and trust stores:** "Trust stores updatable through the management console; signature algorithms follow the firmware." Graded: partial. **Firmness of dates:** firmware 8.0 general availability "planned Q1 2027" in the vendor's written roadmap, no contractual remedy.

Graded: forecast. **ML-KEM and ML-DSA support:** hybrid ML-KEM in 8.0 (forecast); ML-DSA certificates "under evaluation." Graded: unknown. **FIPS 140-3 with PQC in approved mode:** "following module revalidation, date not set." Graded: unknown. **CBOM:** CycloneDX 1.7 provided for firmware 7.4, with the minimum record; accepted. **Non-standardized public-key algorithms:** "none". The statement is attached to the record.

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

- **Scorecard result.** Strategic Blocking; Dimension 3 reads forecast, not the date; the bridging pattern (a PQC-capable proxy in front of the appliance) pre-approved under Activity 7.4; the September 2026 contract renewal carries the crypto-agility clause with the acceptance-test rehearsal.

I.8 Phase 4 – Re-Baseline Protocol Trigger Record

- **Trigger.** The API gateway appliance supplier moved firmware 8.0 general availability from Q1 2027 to Q3 2027 (written notice, January 14, 2027), beyond the bridging capacity recorded in Activity 7.4: the proxy pattern covers the two production gateways and not the four partner-dedicated instances.
- **Review.** Opened at the January SteerCo, January 27, 2027, within the next cycle. Options considered: (a) extend the proxy pattern to the four partner instances, one additional quarter of network engineering; (b) replace the appliance on the partner instances with the challenger product from the champion-challenger evaluation, a different implementation lineage under the Cryptographic Concentration Framework's lineage check; (c) accept a five-month slip on the Wave 3 partner date.
- **Decision.** Option (a), with (b) held as the fallback if the proxy pattern fails interoperability testing with the two largest partners by March 31, 2027. Dates moved: the Wave 3 partner-facing exit moves from April 30 to July 31, 2027. The Tier-1 completion date the board approved as a commitment in the Year-1 plan (December 2027) is unchanged and no regulatory buffer is consumed, so no board notification is required under Activity 0.3. The Wave 3 date was an internal scheduling assumption the board had not been asked to approve. Authority: SteerCo, as the owner of wave completion dates. Minuted; filed to the evidence dossier as RB-01 with the vendor notice attached.

I.9 Phase 3 – Appeals Record under the Ratified Rule

- **Record:** APPEAL-03, decided at the SteerCo of April 22, 2026 under the appeals rule ratified with scoring model v1.0 on March 25, 2026.
- **Entry appealed.** The treasury reporting service (CMDB CI-11207), scored Tier 1 on the March run: Dimension 1 Critical (confidentiality horizon 20 years, over the 15-year line, with partner-facing accessibility High, which meets the first condition of the ordered test), Dimension 4 High (a tier-2 supervisory expectation within five years), Dimension 3 Medium.
- **Appellant and grounds.** The service owner, on evidence rather than on the model: the partner-facing flow scored as the service's exposure was decommissioned on March 31, 2026, and the remaining flows are internal to a network the accessibility factor rates Medium. The confidentiality horizon was taken from a superseded retention schedule and the current classification register gives eight years.
- **Evidence examined.** The March 31 change record and the passive-monitoring observation confirming that the partner endpoint no longer negotiates. The

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

classification register entry dated February 2026 with the information owner's approval.

- **Decision.** Upheld on the accessibility factor and on the horizon: Dimension 1 re-scores Medium, because an eight-year horizon is Medium, internal accessibility is Medium, and the dimension takes the highest of its factors. Under the ordered test of Activity 3.2 the entry leaves Tier 1 (no factor is Critical) and does not meet the Tier 2 condition on Dimension 1 (Medium). It meets the Tier 2 condition on Dimension 4, High on the specific-regulatory-deadline factor, and takes Tier 2 with a start date in the Wave 3 window.

The outcome survives the correction because a second condition catches the entry, and the record shows which one. The model is unchanged. The appeal corrected two inputs and both corrections were written back to the CBOM entry with their evidence. Minuted; filed to the evidence dossier as APPEAL-03. The count of upheld appeals is reported to the SteerCo with the quarterly re-scoring, because a run of upheld appeals on one factor is a model finding rather than a series of exceptions.

I.10 Closure – Decision Record

- **Record:** CLOSE-01, decided by the Board Risk Committee, the level that chartered the program, on March 18, 2029. Scope: the important-service population identified by the inventory snapshot of February 28, 2029, filed with this record.
- **Verification.** Every Tier-1 and Tier-2 service carries a verification disposition: verified within scope (evidence standard of Migration Verification & Program Closure applied; 100% of Tier 1, a 15% sample of Tier 2 per wave with no sampled failure), or an accepted exception with an owner and a re-evaluation date. A disposition on every service is the completeness statement. The migrated count is the first state alone. **Closure report lines.** Unresolved exposed services: eleven Tier-2 systems, all vendor-blocked, each still on a classical path behind a bridging proxy, with the residual exposure recorded from the verification matrix.

Compensating controls: the bridging proxies, verified by hybrid negotiation observed at each proxy in March 2029. Accepted residual risk: the classical path behind each proxy until the vendor firmware ships, accepted by the SteerCo as the risk owner; none of the eleven has an external deadline inside its expiry date, so none is an open obligation. Remediation dates: bridging-pattern expiry dates inside 2029, one per entry, attached.

- **Closure proof.** A rehearsed second algorithm change on one Tier-1 system per track inside the agility window: Track A, X25519MLKEM768 to SecP384r1MLKEM1024 on the customer web edge, 21 h 10 min end to end; Track B, ML-DSA-65 to ML-DSA-87 on the internal code-signing service with verifier acceptance tested on the three production verifier populations, 4 days 6 h end to end. Segment timings filed to the evidence dossier. The agility-debt register shows zero unaccepted debt as a governance condition, with the total reported in full: six entries in all, two Tier-1 (the

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

VPN concentrator awaiting firmware 9.2; a partner-dedicated gateway instance) with funded treatments dated inside 2029 and four Tier-2 with funded treatments in the 2029 plan, none excepted and none unowned. Every Tier-1 and Tier-2 service has one of the four verification dispositions.

- **Decommissioning.** Retired classical roots removed from every trust store in the register; destruction certificates filed for the HSM-held signing keys retired in Waves 2 to 4; decryption keys for the two retained archives kept under the records policy until the archives age out in 2031.
- **Handover.** The QRPM role transfers to the cryptographic governance lead in the CISO function. The Cryptographic Record, continuous discovery, the vendor governance cadence, the SOC detection content, the KRI reporting and the crypto-agility rehearsal cadence transfer to the named owners in the BAU handover register with 2029 operating budgets approved. The evidence dossier is archived to the records system under the ten-year retention the supervisor requires.
- **Decision.** The program is closed within the stated scope with the listed accepted residuals. The residual risk register is accepted by the SteerCo and noted by the committee. The committee records that closure is not a claim that every cryptographic use in the organization is post-quantum, and that the capability the program was chartered to build is now a standing function with an owner and a budget.

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

APPENDIX J: REQUIREMENTS REGISTER

This register is generated from the text of this edition at each build and lists every statement that uses *must*. The four program-level provisions of the Normative Convention come first; every other statement binds inside the activity or section that states it (Normative Convention). A statement that appears here without its activity's conditions is read with them.

Program-level provisions

1. **Phase 5, Deployment Environment Classification.** Organizations must classify their systems by deployment environment before selecting pilot targets or designing hybrid deployments; the classification keys on the module certificate, never on the configuration.
2. **Activity 4.6, gate G1 → G2.** The accepted-gaps register must be signed (Discovery lead, CISO countersign, SteerCo acceptance) before the gate passes.
3. **Activity 4.6, the agility criterion.** At every migration gate from G5 → G6 onward, and at every wave exit, a system counts as migrated only when the first two conditions of the agility criterion hold, which carry the protection result; a system that passes those two and fails the third or the fourth is recorded as migrated with agility debt, and a system that fails either of the first two is not migrated.
4. **Migration Verification & Program Closure.** The program closes only on the closure proof: a rehearsed second algorithm change per track inside the agility window and zero unaccepted agility debt.

Activity-level and section-level statements, in document order

5. *How to Use This Framework, Regulatory Timeline Context.* Readers must verify the status and applicability of every instrument for their own jurisdiction and sector before planning against it.
6. *Phase 0, Purpose.* The mandate must frame the program as a multi-year enterprise transformation.
7. *Phase 0, Purpose.* The mandate must also state two outcomes together.
8. *Phase 0, Purpose, Parallelization note.* Phase 0 is the only phase that must be substantially complete before others begin.
9. *Phase 0, Prerequisites, Organizational prerequisites.* The sponsor must have the organizational standing to secure multi-year budget commitments and cross-functional authority.
10. *Phase 0, Prerequisites, Organizational prerequisites.* If your organization lacks this capability, the governance structure designed in this phase must compensate.

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

11. *Activity 0.3 Establish Governance Structure.* The program must start as a dedicated program with a named QRPM and a steering function that takes the SteerCo's decisions.

12. *Activity 0.3 Establish Governance Structure.* The CFO must be briefed before any board committee sees the funding request, because a multi-year ask that the CFO first meets in the committee papers is contested there.

13. *Activity 0.3 Establish Governance Structure.* 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 exfiltration detection weighted by confidentiality horizon.

14. *Phase 0, Interdependencies, Feeds into.* - Phase 4 – The budget structure and multi-year commitment set the financial constraints within which the roadmap must be built.

15. *Phase 1, Purpose.* Phase 1 must confront this honestly.

16. *Phase 1, Prerequisites, Organizational prerequisites.* 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.

17. *Phase 1, Prerequisites, Organizational prerequisites.* 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.

18. *Activity 1.1 Establish Three Parallel Inventories.* These three inventories must be coordinated through a single governance structure, even if executed by different teams using different tools.

19. *Activity 1.1 Establish Three Parallel Inventories.* For each data category, assign a confidentiality horizon: the number of years the data must remain confidential, derived from, in order of precedence:

20. *Activity 1.2 Deploy Cryptographic Discovery – Layered Approach.* Every discovery finding must carry an evidence grade, in this framework's five tiers:

21. *Activity 1.2 Deploy Cryptographic Discovery – Layered Approach.* Coverage must be measured against an estate enumerated before discovery ran, never against what the tools found.

22. *Activity 1.4 Address the Asset Discovery Problem.* Cloud estates (AWS, Azure, GCP, and other providers) must be queried directly through their APIs and management consoles.

23. *Activity 1.4 Address the Asset Discovery Problem.* 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.

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

24. *Phase 1, Interdependencies, Runs in parallel with.* Staff and budget planning must reflect this.

25. *Phase 1, Common Failures.* The inventory must be maintained in a queryable system (CMDB, dedicated tool, or structured database) with automated refresh.

26. *Phase 2, Prerequisites, Organizational prerequisites.* CycloneDX 1.7 is the format (see Activity 2.1; 1.6 is accepted during a documented transition), but the organization must also decide where the CBOM will be kept: a dedicated tool, a CMDB extension, a version-controlled repository, or a purpose-built database.

27. *Activity 2.1 Select CBOM Format and Tooling.* The table below defines, in this framework's own terms, what a CBOM entry must record, and a CBOM built to it can be expressed as plain CycloneDX.

28. *Activity 2.3 Integrate CBOM into Operational Processes.* - Vendor onboarding: New vendor products must provide CBOM-compatible cryptographic documentation as a procurement requirement.

29. *Activity 2.5 Secure the CBOM and Program Artifacts.* A CBOM must carry a TLP marking set at generation, chosen by the recipients it may be shared with, and is never left at the default.

30. *Phase 2, Interdependencies, Feeds into.* Each Layer 4 gap is a vendor question that Phase 7 must answer.

31. *Phase 2, Common Failures.* The CBOM must be a live, queryable data set integrated into operational processes.

32. *Phase 3, Purpose, Parallelization note.* 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.

33. *Phase 3, Prerequisites, Framework prerequisites (from earlier phases).* 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.

34. *Phase 3, Prerequisites, Framework prerequisites (from earlier phases).* 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.

35. *Activity 3.1 Define Risk Scoring Model.* The SteerCo must ratify the dimensions, the weights, the ordered tier test of Activity 3.2 and an appeals rule before the first scoring run (Activity 0.3).

36. *Activity 3.1 Define Risk Scoring Model.* An entry whose confidentiality horizon or data classification is missing must be left unscored and flagged "classification missing".

37. *Activity 3.1 Define Risk Scoring Model.* It feeds the Crypto-Agility workstream (Activity 0.3) and the Phase 4 roadmap, because a system migrated to PQC without removing its agility debt will be migrated a third time at the next algorithm change,

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

and the roadmap must schedule the debt removal with the migration rather than after it.

38. *Activity 3.1 Define Risk Scoring Model.* The register's burn-down is an agility KPI (Metrics, KPIs & Reporting), and before program closure it must reach zero unaccepted debt: every remaining entry funded with a date or under a SteerCo-signed exception (Migration Verification & Program Closure).

39. *Activity 3.3 Apply Migration Sequencing Logic.* Stores with the longest confidentiality horizons must not fall off the roadmap, and an exposed store (a cloud archive, a backup export, media that has left the organization's control) enters the earliest tranche on its own exposure.

40. *Phase 3, Interdependencies, Runs in parallel with.* The QRA must be re-run at least quarterly as new inventory data arrives, regulatory deadlines shift, vendor timelines change, and NIST standards evolve.

41. *Phase 4, Purpose.* The roadmap must model these external constraints explicitly, not just internal work sequencing.

42. *Phase 4, Prerequisites, Framework prerequisites (from earlier phases).* The Phase 0 budget structure, including the phased funding model and benefit tracking approach – provides the financial framework within which the roadmap must fit.

43. *Phase 4, Prerequisites, Framework prerequisites (from earlier phases).* The roadmap must include realistic effort estimates, which require at least a rough understanding of what needs to change in each system.

44. *Activity 4.2 Structure the Multi-Year Roadmap.* The roadmap must show the Crypto-Agility workstream's milestones on the same plan as the migration waves.

45. *Activity 4.4 Establish PMO Structure for Scale.* Critical dependency chains (must be mapped and tracked):

46. *Activity 4.4 Establish PMO Structure for Scale.* Upstream Activity | Must Complete Before | Why

47. *Activity 4.4 Establish PMO Structure for Scale.* Risk scoring (Phase 3) | Pilot target selection (Phase 5) | Pilots must target the highest-priority systems, not the most convenient ones.

48. *Activity 4.4 Establish PMO Structure for Scale.* The vendor dependency is usually the longest segment on the critical path, which is why vendor engagement must start early, before the full CBOM is complete.

49. *Activity 4.6 Define Milestone Gates.* The gate must not pass it as migrated.

50. *Activity 4.8 Re-Baseline Protocol.* Any one of the following must open a re-baseline review within the next SteerCo cycle: actual effort on a wave exceeding the estimate by a factor the SteerCo set in advance (1.5 is a workable default); a Strategic Blocking vendor slipping a committed date beyond the bridging capacity recorded in Activity 7.4; a pilot failure of a class the pilot plan defined as blocking (an interoperability failure with a Tier-1 counterparty, a security finding in the

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

deployed implementation); or movement of an external deadline the roadmap was built against.

51. *Phase 4, Interdependencies, Feeds into.* - 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.

52. *Phase 4, Interdependencies, Runs in parallel with.* The roadmap review process must accommodate these inputs.

53. *Phase 4, Common Failures.* The SteerCo must have decision authority over budget, timelines, risk acceptance, and vendor escalation.

54. *Phase 5, Purpose, Parallelization note.* 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.

55. *Phase 5, Prerequisites, Organizational prerequisites.* Pilots must be tested in environments that include the same middleboxes, load balancers, WAFs, monitoring tools, and partner endpoints as production.

56. *Phase 5, Prerequisites, Organizational prerequisites.* 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.

57. *Phase 5, Prerequisites, Organizational prerequisites.* The underlying cryptographic libraries (OpenSSL, BoringSSL, Bouncy Castle, platform-native crypto APIs) must support the target PQC algorithms at the required versions.

58. *Phase 5, Deployment Environment Classification, Understanding Deployment Constraints.* Before selecting pilot targets or designing hybrid deployments, organizations must classify their systems by deployment environment, because that determines when PQC can enter production, regardless of technical readiness.

59. *Phase 5, Deployment Environment Classification, Understanding Deployment Constraints.* A vendor statement that a module is "hybrid-validated" must be read against that definition.

60. *Phase 5, Deployment Environment Classification, Understanding Deployment Constraints.* CNSA 2.0 / National Security System environments must comply with NSA CNSA 2.0 timelines.

61. *Phase 5, Deployment Environment Classification, Understanding Deployment Constraints.* Under the 2024 update to CNSSP 15, new NSS acquisitions must be CNSA 2.0 compliant from January 1, 2027; equipment and services that cannot support CNSA 2.0 are phased out by December 31, 2030; and CNSA 2.0 algorithms are mandated for use by December 31, 2031, each unless otherwise noted in the applicable profile.

62. *Phase 5, Two-Track Migration Model, Track A (Confidentiality) and Track B (Integrity and Authentication).* A deployment that does not meet section 8 of SP 800-

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

208 is not a conforming deployment, and the verifiers must accept the parameter set chosen.

63. *Phase 5, Two-Track Migration Model, Track A (Confidentiality) and Track B (Integrity and Authentication).* Its signing depends on floating-point arithmetic that must be implemented in constant time, which creates side-channel surface that validated-module support has been slow to cover, so FN-DSA is the constrained-environment option after standardization and validation, not the HSM default.

64. *Activity 5.1 Select Pilot Targets.* - Rollback capability – must be able to revert to pre-pilot state if issues arise

65. *Activity 5.2 Design Hybrid Deployments.* The classical component is non-approved in X25519MLKEM768 and additionally approved in the NIST-curve groups, which matters only where classical-only operation must also remain approved).

66. *Activity 5.2 Design Hybrid Deployments.* The guidance is unchanged: hybrid X25519MLKEM768 is the default wherever confidentiality must outlive the connection; pure ML-KEM is a compliance profile for regimes that prescribe it, adopted once the RFC has published (Appendix D).

67. *Activity 5.2 Design Hybrid Deployments.* Pilot design for long-lived channels (database replication links, message-queue connections, VPN tunnels, service-mesh sidecars) must force a fresh handshake on a defined interval after hybrid enablement, must require the hybrid group on resumption and disable 0-RTT for the flows in scope, and may expire pre-migration tickets as the simpler policy.

68. *Activity 5.2 Design Hybrid Deployments.* Programs that piloted in 2023 and 2024 on draft identifiers (the X25519Kyber768Draft00 and SecP256r1Kyber768Draft00 code points, and the draft hybrid names used in libraries, load balancers and proxies of that period) must retire them on a dated schedule: RFC 10024 obsoletes the pre-standard Kyber768 code points and marks them Recommended=D in the IANA registry.

69. *Activity 5.2 Design Hybrid Deployments.* Every wave must include an algorithm-change rehearsal in staging on at least one system in the wave: a parameter-set change (ML-KEM-768 to ML-KEM-1024 on a key-establishment path), a hybrid-to-composite switch on a signing path, or a certificate-signature-algorithm change on an internal issuer.

70. *Activity 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.

71. *Activity 5.6 Decide the Data-at-Rest Strategy.* 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.

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

72. *Phase 5, Common Failures.* A pilot that succeeds on a developer sandbox proves little; pilots must prove that PQC works on the highest-priority, highest-risk systems.

73. *Phase 5, Common Failures.* When a pilot causes unexpected failures, (and some will) the team must be able to revert within minutes, not hours or days.

74. *Phase 6, Purpose, Parallelization note.* 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.

75. *Phase 6, Prerequisites, Framework prerequisites (from earlier phases).* These must be planned and funded in the Phase 4 roadmap.

76. *Phase 6, Prerequisites, Organizational prerequisites.* This work must be done carefully.

77. *Activity 6.1 PKI Modernization.* Publicly trusted TLS subscriber certificates must assert serverAuth only.

78. *Activity 6.2 HSM and KMS Modernization.* - HSM and KMS key generation, and any software module that generates long-lived PQC keys, must draw from an entropy source validated under SP 800-90B.

79. *Activity 6.3 Network Infrastructure Assessment.* PQC impacts network infrastructure in several ways that must be tested before production deployment:

80. *Phase 6, Interdependencies, Runs in parallel with.* When a vendor misses a delivery milestone, Phase 6 must invoke a bridging strategy rather than simply delaying.

81. *Phase 6, Common Failures.* Every network path, middlebox, load balancer, and TLS termination point must be tested, not just the application endpoints.

82. *Phase 7, Purpose.* 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.

83. *Phase 7, Prerequisites, Organizational prerequisites.* The executive sponsor must be prepared to escalate to vendor C-suite leadership, invoke contractual leverage, and if necessary authorize evaluation of alternative products.

84. *Phase 7, Prerequisites, Organizational prerequisites.* The framework provides a classification matrix and questionnaire structure (Activities 7.1–7.2), but the organization must designate staff to operate this process.

85. *Phase 7, Prerequisites, Organizational prerequisites.* The vendor governance process must include open-source dependency tracking with the same rigor applied to commercial vendors.

86. *Activity 7.3 Insert PQC Requirements into Procurement.* It must be tested rather than warranted: acceptance testing before go-live includes a rehearsed algorithm

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

change on the product, executed by the customer's team using the rehearsal pattern in Activity 5.2, with the elapsed times recorded and the rollback exercised.

87. *Phase 7, Interdependencies, Feeds into.* The roadmap must model these vendor dependencies explicitly so that timeline changes propagate to affected milestones automatically.

88. *Migration Verification & Program Closure, Verification Evidence Standard.* Passive monitoring (the Phase 1 continuous discovery infrastructure) must show the system negotiating the expected hybrid or PQC algorithms in production traffic.

89. *Migration Verification & Program Closure, Verification Evidence Standard.* - Negative testing where classical-only must be refused.

90. *Migration Verification & Program Closure, Decommissioning Classical Material.* Public verification material and historical validation evidence are preserved where signed records must remain verifiable (Appendix H), and the classical component of an approved hybrid deployment stays active by design.

91. *The Seven Program Foundations, Metrics, KPIs & Reporting, Board-Level KPI Pack (Report Quarterly).* Agility-debt burn-down | Entries closed on the Phase 3 agility-debt register, per quarter | Any quarter with no reduction triggers workstream review; before closure every remaining entry must carry a funded treatment with a date or a SteerCo-signed exception (zero unaccepted debt)

92. *The Seven Program Foundations, Regulatory & Standards Alignment Map, Deadline Archetypes.* Procure-by | New acquisitions must carry PQC support from a date | CNSA 2.0: new National Security System acquisitions CNSA 2.0-compliant from 2027 | Tier 1, a procurement gate with a named enforcer

93. *The Seven Program Foundations, Regulatory & Standards Alignment Map, Algorithm Sovereignty and Standards Fragmentation.* For a multinational the question is no longer hybrid posture but algorithm portfolio: which PQC families the architecture must support, selected per jurisdiction.

94. *The Seven Program Foundations, Regulatory & Standards Alignment Map, Algorithm Sovereignty and Standards Fragmentation.* 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.

95. *The Seven Program Foundations, Skills & Team Structure, Sustaining Capability Beyond the Migration.* The skills and structures built for PQC migration must outlast the program itself.

96. *The Seven Program Foundations, SOC Implementation.* The Security Operations Center must develop detection capabilities for quantum-related threats, build a threat intelligence function that tracks CRQC developments and PQC implementation

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

vulnerabilities, and maintain incident response playbooks for scenarios that have no equivalent in the SOC's current library.

97. *The Seven Program Foundations, SOC Implementation.* This section specifies what the SOC must build, in what order, and with what resources.

98. *The Seven Program Foundations, SOC Implementation, Prerequisites.* The SOC must have access to a queryable, continuously updated cryptographic posture registry that maps systems to their expected cryptographic configuration.

99. *The Seven Program Foundations, SOC Implementation, Prerequisites.* For each system in scope, the registry must specify three things:

100. *The Seven Program Foundations, SOC Implementation, Prerequisites.* The registry must be machine-readable and integrated with the SIEM, ideally through an API or automated feed that updates as migration progresses.

101. *The Seven Program Foundations, SOC Implementation, Detection Use Cases.* 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.

102. *The Seven Program Foundations, SOC Implementation, Detection Use Cases.* Detection rules must account for these fragmentation-induced retries, and the migration program (Phase 6) should flag middlebox upgrades as a prerequisite for reliable hybrid deployment.

103. *The Seven Program Foundations, SOC Implementation, Incident Response Playbooks.* The SOC must be briefed on the rotation schedule, with suppression rules for expected activity, while maintaining detection for actual adversary activity.

104. *The Seven Program Foundations, GRC Implementation.* 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.

105. *The Seven Program Foundations, GRC Implementation, Risk Appetite and Tolerance.* 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.

106. *The Seven Program Foundations, GRC Implementation, Risk Appetite and Tolerance.* "All newly deployed systems from Q3 2026 onward must implement crypto-agile architecture as a mandatory design requirement.

107. *The Seven Program Foundations, GRC Implementation, Qualified Trust Services and Electronic Signature Law.* Some archives have a legal effect that must outlive the classical algorithms protecting them.

108. *The Seven Program Foundations, GRC Implementation, The GRC-SOC Handoff.* The exception and deferral register, which the SOC must ingest as a dynamic suppression list.

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

109. *Sector Adaptation Notes, Critical Infrastructure / OT.* - Safety constraints: Any change to OT cryptography must go through safety case re-certification in many jurisdictions

110. *Sector Adaptation Notes, Government & Defense.* Key generation and signing must enter production through a module whose certificate lists the operation as an approved service, and SP 800-208's conditions for stateful hash-based signing keys apply unchanged.

111. *Appendices, Appendix A: Algorithm Quick Reference.* Compact lattice-based signatures suitable for bandwidth-constrained environments, though signing requires floating-point arithmetic that must be implemented in constant time.

112. *Appendices, Appendix C: Mosca's Inequality – The Decision Framework.* If X + Y > Z, you must act now.

113. *Appendices, Appendix C: Mosca's Inequality – The Decision Framework.* - X = Security shelf-life: the number of years the data you are protecting must remain confidential

114. *Appendices, Appendix D: Hybrid Approach Jurisdictional Compliance Matrix.* NSA CNSA 2.0 | Accepted as interim only | Hybrid acceptable during transition; end-state must be CNSA 2.0 compliant (pure PQC)

115. *Appendices, Appendix F: Common Objections and Evidence-Based Responses.* Mosca's Inequality (X + Y > Z, where X is data sensitivity lifetime, Y migration duration and Z time until CRQC) shows that organizations with long-lived sensitive data and multi-year migration timelines must start now regardless of the CRQC date.

116. *Appendices, Appendix F: Common Objections and Evidence-Based Responses.* Many critical systems depend on multiple vendors, all of whom must ship PQC support before migration can proceed.

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

ABOUT THIS VERSION

This framework is versioned and updated as standards evolve, vendor products mature, and regulatory deadlines shift. Version 3.0 is a major version. It follows the June 2026 v2.1, a targeted update to the June 2026 v2.0 release, which updated the March 2026 v1.1.

WHAT'S NEW IN V3.0

Version 3.0 is a major version because crypto-agility becomes the program's organizing method and its measured end state. The eight phases, the two-track model, the PKI architecture fork, the deployment environment classification and every v2.x position not corrected here stand unchanged. The mandate is readiness against dated obligations; the method is crypto-agility; the first exercise is the PQC migration; the end state is an organization that changes its cryptography on a planned date.

Crypto-agility as method and end state, and the register. The charter states readiness and agility together (Phase 0); crypto-agility is a funded workstream with an owner, and Discovery is separated from the permanent Cryptographic Record (Activity 0.3); the milestone gates and wave exits gain an agility criterion (Activities 4.6 and 5.4); every wave carries an algorithm-change rehearsal, and closure requires a rehearsed second change on Tier-1 systems inside the agility window and zero unaccepted agility debt (Activity 5.2; Migration Verification & Program Closure); Phase 3 produces an agility-debt register; the Crypto-Agility foundation gains a mechanisms catalogue, a rehearsal-based measurement method and agility KPIs with thresholds. Normative convention added (must, should, may). Every phase gains an Applying This Phase block, a closing common-failure line, and challenge questions for the SteerCo. Appendix I provides worked artifacts across the phases and at closure.

Sibling properties and the regulatory restructure. Phase 2 defines this framework's own minimum CBOM record and cites the Applied Quantum CBOM Profile by minimum version as its machine-readable carriage, never as a requirement; the Cryptographic Concentration Framework is cross-referenced from Phases 3 and 7, the KRI cascade and Appendix G, and no concentration figure is computed here. Discovery findings carry evidence grades, and coverage is measured against an enumerated denominator with a signed accepted-gaps register. Regulatory content states the mechanism (deadline archetypes, enforcement tiers, the regulatory intelligence process) and three anchor instruments; the jurisdiction roster is printed as a dated snapshot and maintained between editions on the Global PQC Migration Clock page, updated from the PQC

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

Migration Brief, which is now the framework's interim-update channel. Commercial vendor and product names moved to the front-matter Tool and Vendor References with an as-of date; open-source libraries and kits stay named in the body.

Two loud corrections. HNDL is stated as inference from demonstrated capability rather than as observed harvesting (Two-Track Migration Model, Appendix F), and Appendix F quotes three published statements, from CISA, NSA and NIST, from AIVD, CWI and TNO, and from Federal Reserve Board staff, each dated and linked; SOC Use Case 5 is reframed as exfiltration detection weighted by confidentiality horizon, and the "HNDL Exposure Score" is retired.

Method decisions adopted after review. Nine decisions taken on the external reviews of the draft: the scoring arithmetic stated once, with the highest factor scoring its dimension, inapplicable dimensions excluded and weights renormalized, and unknown values unscored until resolved (Activity 3.2); a delivery state on every scored entry, so that a blocked migration keeps its risk tier and appears on the roadmap against its unblocking date (Activities 3.2 and 4.2); two results in every gate and wave-exit record, protection and agility, read separately (Activities 4.6 and 5.4); closure at zero unaccepted agility debt, with a verification disposition for every Tier-1 and Tier-2 service (Migration Verification & Program Closure); rollout, wave-exit and closure gates named as different decisions (Activity 4.6); the evidence grade attached to a claim rather than to an inventory row (Activity 1.2); the three lines of defense stated by function rather than by title (Activity 0.3; GRC Implementation); one maturity scale across the phases and the enterprise model; a requirements register generated from the text (Appendix J); and Marin's Law labeled a design heuristic (Crypto-Agility foundation).

Four decisions on the external review of September 13, 2026. The tier is assigned by an ordered qualitative test and the composite score orders entries inside a tier, with no score band published (Activity 3.2; Appendix I.3); a verification matrix states positive evidence, the negative or boundary test and the residual exposure for each cryptographic function the negotiated-protocol test does not reach, and the agility criterion reads it for those functions (Migration Verification & Program Closure; Activities 4.6 and 5.4); the minimum record gains the five fields scoring consumes, marked recorded at Phase 3, and the unit of a row is settled as the cryptographic use (Activity 2.1); and the HSM partition failover is recorded as a failover and state-management exercise, not as an algorithm-change rehearsal (Activity 5.2; Crypto-Agility foundation).

Governance instruments and standards currency. Scoring-model ratification with an appeals rule and model-risk treatment; the Re-Baseline Protocol (Activity 4.8); verifier updatability as a Dimension 2 factor; discovery access governance; the classical-findings reporting rule; the CFO pre-brief; semi-annual maturity re-assessment with trend reporting; the certificate-authority and non-standardized-algorithm questions in the supplier questionnaire, with firmness grading of vendor dates and the provider-

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

roadmap reading rule; a front-end stage for unverified cryptanalytic claims in Playbook 1 with a grading standard in Strategic CTI; qualified trust services and electronic signature law in GRC Implementation; key generation and entropy in Phase 6; the client-authentication sunset and signature placement in mixed chains in PKI modernization; first-layer selection by architecture pattern, long-lived session and resumption handling, and pre-standard code-point retirement in Phase 5. Standards currency: RFC 9954 and RFC 10024 for hybrid TLS, with the registry rule for pure ML-KEM; RFC 9881, RFC 9882 and RFC 9814 for S/MIME; RFC 9980 for OpenPGP; the 2026 FIPS 140-3 Implementation Guidance updates; CycloneDX 1.7 and ECMA-424; FN-DSA and HQC status; the federal CBOM guidance mandate, phrased as directed rulemaking. Payments and Digital Assets join the Sector Adaptation Notes. All six sector extensions are updated to v3.0 and publish with this edition.

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 deployment-status 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. 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. *Corrected at v3.0:* that statement was

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

false when v2.0 published. CMVP certificate #5247, a software module validated on April 27, 2026, lists ML-KEM key generation and encapsulation-decapsulation at ML-KEM-768 and ML-KEM-1024 among its approved algorithms; the v3.0 Deployment Environment Classification (Phase 5) replaces the calendar-based gap with a per-product validation decision.

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, superseded at v3.0 by exfiltration detection weighted by confidentiality horizon) 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 39 crypto-agility alignment; algorithm pipeline updates (FN-DSA, HQC, 9 additional signature candidates); current PQC deployment status data; a multinational

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

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. Movement between editions is covered in the PQC Migration Brief (pqcmigrationbrief.com) and, for jurisdiction deadlines, on the Global PQC Migration Clock page.

STANDARDS IN PROGRESS

NIST IR 8547 (Initial Public Draft November 2024; still a draft at the date of this edition) 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 founder and CEO of Applied Quantum, and the founder of Quantum Academy (QuantumAcademy.com). PostQuantum.com is his personal blog. 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 standards, products and deadlines change, and the book supplies the depth that a methodology document omits by design.

QUANTUM ACADEMY

Quantum Academy (QuantumAcademy.com), founded by the author, offers trainings and certifications on the topics this framework covers, from executive education on quantum risk to practitioner courses on PQC migration, crypto-agility and cryptographic inventory, for the roles and training levels described in Skills & Team Structure.

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

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 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.