electronics

MDPI

Article

A Quantum-Adjusted Risk Model for Enterprise Infrastructure Across Data In Transit, In Use, and At Rest

Simas Krušniauskas, Šarūnas Grigaliūnas *, Rasa Brūzgienė and Mert Cayir

Department of Computer Sciences, Kaunas University of Technology, LT-51368 Kaunas, Lithuania; simas.krusniauskas@ktu.lt (S.K.); rasa.bruzgiene@ktu.lt (R.B.); mert.cayir@ktu.edu (M.C.)
* Correspondence: sarunas.grigaliunas@ktu.lt

Abstract

Enterprise infrastructure operators face a critical challenge in prioritizing post-quantum migration, as quantum-related risk is not uniformly distributed across data in transit, in use, and at rest. Existing assessments rely on system-level evaluations or protocol-specific analyses, which do not capture the heterogeneity of exposure across infrastructure layers. This paper extends the Quantum-Adjusted Risk Scoring (QARS) model introduced in into an evidence-based, layer-specific framework that evaluates in-transit, in-use, and at-rest data separately. QARS applies a unified five-factor scoring framework separately to each data state and introduces a quantum-vulnerability attenuation mechanism grounded in Grover-bounded residual security that prevents overstating urgency for non-Shor-vulnerable symmetric protection. Observable host-level evidence determines the binary and ratio descriptors used by the model, while the fixed affine mapping coefficients are treated as transparent semi-quantitative calibration parameters. These coefficients are documented separately and subjected to coefficient-level sensitivity analysis to evaluate whether the reported layer ordering depends on their nominal values. The model is demonstrated through an illustrative controlled experiment using real infrastructure observations. Strengthening storage protection reduces the aggregate system risk from 0.707 (high) to 0.414 (moderate), a 41.5% reduction. However, the maximum-layer score remains high (0.657), indicating that the transport layer continues to dominate migration urgency. Sensitivity analysis confirms that the dominance of the transport layer is stable under wide perturbations of the calibration parameters. These findings demonstrate that risk reduction in one layer does not eliminate overall exposure but shifts the dominant vulnerability. By distinguishing between overall system posture and the most critical remediation priority, QARS supports infrastructure operators in identifying high-risk components and planning structured, evidence-based post-quantum migration.

Keywords: post-quantum cryptography; quantum risk assessment; critical infrastructure security; evidence-based scoring; encrypted storage; risk prioritization

Academic Editor: Hung-Yu Chien

Received: 15 April 2026
Revised: 22 May 2026
Accepted: 4 June 2026
Published: 9 June 2026

1. Introduction

The prospect of large-scale quantum computing has transformed post-quantum cryptography from a long-range research concern to an immediate infrastructure planning problem. Shor's polynomial-time algorithms for integer factorization and discrete logarithms [1] render every widely deployed asymmetric primitive based on RSA, finite-field Diffie–Hellman, and elliptic-curve cryptography insecure once a cryptographically relevant

quantum computer (CRQC) becomes available, while Grover's quantum search algorithm [2] delivers only a square-root speed-up against unstructured symmetric primitives, leaving AES-256 and AES-128-XTS with effective post-quantum security levels of approximately 128 and 64 bits per block, respectively. This asymmetry between Shor-vulnerable public-key mechanisms and Grover-bounded symmetric primitives is the cryptographic basis for differentiated migration urgency across data layers, but it is rarely reflected in current organizational risk-assessment practice.

The shift toward concrete migration planning became more tangible with the publication of NIST's first post-quantum standards, such as FIPS 203 for ML-KEM, FIPS 204 for ML-DSA, and FIPS 205 for SLH-DSA, which established a deployable baseline for systems that currently depend on public-key primitives vulnerable to quantum attack [3–5]. The wider NIST PQC project provides a broader standardization context [6]. Complementary national and international guidance now extends these primitives into operational migration playbooks: the NCCoE practice guide on Migration to Post-Quantum Cryptography (NIST SP 1800-38) [7], the ETSI repeatable framework for quantum-safe migrations TR 104 016 [8], the prior ETSI staged-migration recommendations TR 103 619 [9], the NSA Commercial National Security Algorithm Suite 2.0 [10], NIST SP 800-131A on cryptographic transitions [11], and the UK NCSC migration timelines [12] all converge on the same conclusion: cryptographic discovery and prioritization, rather than algorithm selection, is the critical bottleneck for organizational readiness [13]. Yet, the existence of standardized algorithms and structured migration frameworks does not, by itself, tell operators where migration should begin in a heterogeneous estate. Enterprise and national infrastructure environments are rarely cryptographically uniform [14,15]: some services expose transport channels to public networks, some process short-lived operational state in memory, and others retain data whose confidentiality must be preserved for many years. Consequently, post-quantum exposure should be assessed separately across transport, runtime, and storage contexts because these contexts rely on distinct protocol, in-use protection, and storage protection mechanisms [16,17].

Existing work in this area can be understood in the following broad strands. The first focuses on algorithm standardization and deployment readiness [6,13], establishing the cryptographic primitives and migration sequencing that can support a post-quantum transition. The second concerns secure communication protocols, especially TLS and SSH, which remain foundational to transport security in enterprise environments [18–20]. The third addresses platform- and system-level protection of operational and stored data; for example, Redis documentation covers TLS-based in-transit protection [21] and persistence/at-rest behavior [22]. Post-quantum transition studies examine their overhead, authentication, hybridization, and deployment feasibility [23,24]. A fourth, organizationally oriented strand has emerged around crypto-agility and bill-of-materials approaches, including CARAF as a 5D risk assessment framework [25], a crypto-agility maturity model [26], and OWASP CycloneDX Cryptography Bill of Materials (CBOM) for cryptographic asset inventory [27]. Together, these strands provide the technical and procedural foundations for migration. However, they do not fully resolve a practical assessment gap: organizations still lack a concise, evidence-based method for comparing post-quantum exposure across different data states within the same infrastructure. CARAF [25] provides a qualitative ordinal risk model but does not separate exposure by data state; ETSI TR 104 016 [8] prescribes a procedural framework but does not produce a layer-resolved numerical score; and NIST SP 1800-38 [7] concentrates on cryptographic discovery and interoperability rather than risk quantification.

This gap is not merely theoretical. In operational practice, migration readiness is often judged at the level of an application, host, or service inventory [13,14]. Such coarse

assessment can obscure important differences between a publicly reachable legacy transport surface, a tightly controlled in-memory processing component, and a storage layer protected by strong symmetric encryption. Treating all of these layers as if they carry the same quantum urgency may lead to distorted priorities. Exposed communication channels that still depend on legacy public-key mechanisms can create urgent interception and delayed-decryption risk–also known as "harvest now, decrypt later" attacks [28]–whereas at-rest data protected by contemporary symmetric controls should not necessarily be assigned the same level of post-quantum urgency because Grover's algorithm degrades but does not break appropriately keyed AES [2]. Confidential-computing mechanisms provide a separate class of in-use isolation [23,29]. The central problem addressed in these papers is therefore how to construct a risk model that captures these differences without losing the ability to compare them at the system level.

This problem is especially important for enterprise and critical infrastructure operators, where service continuity, regulatory obligations under regimes such as the EU NIS2 Directive [30] and the EU Cyber Resilience Act, and limited modernization resources require defensible prioritization [31–34]. A model that overstates the urgency of already-protected layers may divert effort away from more vulnerable components, while a model that hides transport exposure behind aggregate system labels may delay the remediation of the true bottleneck. For post-quantum migration planning to be operationally useful, it must distinguish between total system posture and the most urgent unresolved layer, and it must ground its conclusions in observable infrastructure evidence rather than abstract assumptions.

To address this need, this paper extends the Quantum-Adjusted Risk Scoring (QARS) model originally introduced in [35]–which formalizes post-quantum risk along the timeline, sensitivity, and exposure dimensions–into a layer-resolved, evidence-based framework spanning three principal data states: in transit, in use, and at rest. The novelty of this work lies in the following:

1. The proposed method evaluates each data state layer separately while retaining a common scoring logic, allowing transport, processing, and storage to be compared without collapsing their distinct security properties.
2. The model derives its parameters from concrete host evidence, including exposed services, legacy protocol conditions, weak-cipher observations, runtime controls, proxy wrapping, and disk encryption status, rather than from synthetic laboratory assumptions.
3. This study introduces a quantum-vulnerability attenuation term, grounded in Grover-bounded residual security, for cases in which the dominant protection is symmetric and therefore not directly susceptible to Shor-type compromise. This avoids overstating the post-quantum urgency of protected storage relative to legacy public-key communication channels.
4. This paper reports both a weighted aggregate score and a maximum-layer score, thereby separating overall risk posture from immediate remediation priority.
5. The model's calibration parameters (T_Q, T_B, attenuation coefficients α and β, aggregation weights, and risk thresholds) are subjected to one-at-a-time and multi-dimensional sensitivity analysis, and the results are compared side-by-side with CARAF [25], ETSI TR 104 016 [8], NIST SP 1800-38 [7], the Mosca timeline argument [28], and CycloneDX CBOM [27].

This research is guided by the following research questions:

- RQ1: Does a layer-specific, evidence-based scoring model provide a more meaningful characterization of post-quantum exposure than a uniform system-level assessment?

- RQ2: How does the system-level score change when the at-rest layer is hardened with LUKS2-backed symmetric protection while the in-transit and in-use layers remain unchanged?
- RQ3: Does such hardening change the dominant migration priority, or does transport exposure remain the principal bottleneck?

These questions connect cryptographic modernization to operational decision-making. They help determine not only whether a control reduces risk but also where limited remediation efforts should be directed first.

Accordingly, this paper extends the canonical quantum-adjusted scoring formulation associated with the Mosca timeline argument [28,36] and the CARAF threat-vector decomposition [25] by incorporating an attenuation mechanism for cases in which symmetric protection constitutes the dominant security control. It also establishes an auditable mapping between host-level observations and layer-specific parameters, including confidentiality duration, migration effort, sensitivity, and exposure. The proposed solution is demonstrated through an illustrative case study spanning transport, processing, and storage conditions, and its parameter sensitivity is examined across the calibration ranges discussed in Section 6.5. Statistical generalization to a production estate is identified as a limitation and a direction for future work within enterprise infrastructure. By reporting both aggregate and maximum-layer scores, this work provides complementary perspectives that enhance strategic risk communication while supporting the prioritization of near-term remediation actions.

The remainder of this paper is organized as follows. Section 2 reviews related work on post-quantum migration, infrastructure security assessment, and protection mechanisms across data states. Section 3 presents the proposed Quantum-Adjusted Risk Scoring (QARS) methodology, including the canonical scoring model, the quantum-vulnerability attenuation mechanism, the evidence-to-feature mapping process, and the system-level aggregation logic; the differentiation of QARS from prior frameworks; the calibration of central parameters; and the treatment of cross-layer interdependencies. Section 4 describes the source data and the host-level evidence extracted for the experimental evaluation, including the treatment of measurement uncertainty. Section 5 explains the experimental design, the compared scenarios, and the reproducibility artifacts. Section 6 reports the evaluation results, including layer-level scores, system-level scores, sensitivity analysis, and a comparative analysis against existing frameworks. Section 7 discusses the implications of the results for post-quantum migration planning in enterprise infrastructure, including an operational transport-layer modernization plan and the ceiling effect of the clip operator. Finally, Section 8 concludes this paper and outlines directions for future work.

2. Related Works

Research relevant to quantum-adjusted infrastructure risk assessment spans several interconnected domains, including post-quantum migration management, protocol-level transition to quantum-safe communication, critical infrastructure cyber-risk modeling, and mechanisms for protecting data across their lifecycle. While each of these areas has advanced independently, the literature does not yet provide a unified, evidence-based framework for evaluating post-quantum exposure across data in transit, in use, and at rest within operational environments (Table 1).

Table 1. Comprehensive mapping of related work to data layers and risk-relevant factors.

Research Area	References	Data Layer	Key Factors Addressed	Limitations for QARS
Post-quantum migration and crypto-agility	[14,15,26,28,37,38]	Cross-layer	Migration effort (T_{mig}), governance, crypto-agility maturity, dependency tracking, lifecycle management, system coordination	No quantitative scoring; does not separate transit, in-use, and at-rest exposure
Protocol-level transition (TLS/SSH)	[18–20,39–44]	In-transit	Public exposure (E), protocol legacy, handshake overhead, hybrid authentication, interoperability, deployment feasibility	Focused only on transport; no cross-layer comparison
Critical infrastructure risk and resilience	[31–34]	Cross-layer	Sensitivity (S), cascading failures, interdependencies, resilience, operational impact, system-wide consequences	Not post-quantum specific; lacks cryptographic-layer differentiation
Runtime protection and confidential computing	[23,24,29,45]	In-use	Runtime exposure (E), execution isolation, trusted environments, access control, secure computation, data-in-use protection	Does not integrate with transport or long-term storage risk
Storage security and data lifecycle protection	[46]	At-rest	Confidentiality duration (T_{conf}), persistence, secure storage, deletion guarantees, lifecycle protection	No post-quantum differentiation; no system-level comparison

A first line of work focuses on post-quantum migration planning and crypto-agility. Nethen et al. [14] proposed a structured migration management process for transitioning from classical to post-quantum cryptography, emphasizing governance, asset inventory, and phased deployment. Similarly, Malina et al. [15] examined the deployment of quantum-resistant mechanisms in intelligent infrastructures, highlighting challenges related to heterogeneity, interoperability, and operational continuity. Complementary work on crypto-agility maturity models further emphasizes that migration readiness depends not only on technical capabilities but also on organizational processes and lifecycle management [26,37]. Additional studies highlight the importance of dependency tracking and system-wide coordination when migrating cryptographic infrastructures [28,38]. These studies demonstrate that migration is a system-level problem requiring coordination across protocols, platforms, and services, but they do not provide quantitative mechanisms for comparing risk across infrastructure layers.

A second research stream investigates protocol-level implications of post-quantum cryptography, particularly for TLS and SSH. Early work quantified the overhead introduced by post-quantum primitives in secure communication protocols [18], while subsequent studies evaluated performance and feasibility in modern implementations [19], embedded systems [39], and resource-constrained environments [40]. Other contributions explored transition mechanisms such as hybrid authentication through mixed certificate chains [41] and techniques to reduce certificate overhead in post-quantum TLS [42]. TPM-assisted approaches have further demonstrated the feasibility of integrating post-quantum authentication into hardware-supported environments [43]. Additional studies emphasize interoperability challenges and the need for hybrid cryptographic schemes during transition phases [20,44]. Although this body of work provides detailed insights into performance and deployment trade-offs, it typically focuses on individual protocols rather than system-wide exposure across multiple data states.

A third body of literature addresses cyber-risk and resilience in critical infrastructure. Carvalho et al. [31] introduced an impact assessment framework that models the cascading effects of cyberattacks on critical systems, while Segovia-Ferreira et al. [32] surveyed resilience strategies for cyber-physical systems, emphasizing the need for metrics that support preparedness, absorption, and recovery. Additional research highlights cascading failure effects and interdependency modeling as central to understanding systemic cyber-risk [33]. Studies on interdisciplinary risk assessment further demonstrate that infrastructure risk must incorporate operational, technical, and organizational dimensions simultaneously [34]. These approaches extend risk analysis beyond isolated vulnerabilities toward system-level consequences, but they are not tailored to post-quantum threats and do not explicitly distinguish between transport, runtime, and storage exposure.

A fourth research direction focuses on the protection of sensitive data during processing and storage. Recent work on confidential computing has explored the trade-offs between cryptographic protection and trusted execution environments for securing data in use [23]. Systematization studies have further demonstrated how confidential computing can support secure data processing and machine learning workloads [24,45]. Enclave-based architectures and secure execution environments provide controlled access to sensitive datasets while reducing exposure during computation [29]. In parallel, earlier foundational work on secure storage and deletion highlights the importance of long-term confidentiality guarantees and data lifecycle control [46]. These studies are particularly relevant to the in-use and at-rest dimensions of infrastructure security, as they highlight the importance of isolation, access control, and encryption strength. However, they do not address how such protections should be compared against transport-layer vulnerabilities or long-term storage risks in a post-quantum setting.

The existing literature provides important foundations for post-quantum migration, protocol engineering, infrastructure resilience, and data protection. Nevertheless, a key gap remains: there is no concise, evidence-based framework that (i) evaluates post-quantum exposure separately across data states, (ii) derives scoring inputs from observable host conditions, and (iii) enables direct comparison between layers while accounting for differences in cryptographic vulnerability. The present work addresses this gap by introducing a layer-specific, evidence-driven scoring model that integrates these perspectives into a unified risk assessment approach.

3. Methodology for Quantum-Adjusted Risk Scoring Model

This section presents the proposed Quantum-Adjusted Risk Scoring (QARS) model for evaluating post-quantum exposure across enterprise infrastructure. The model is designed to capture the asymmetric nature of quantum-related risk across three data states: in transit, in use, and at rest. The overall workflow of the model is illustrated in Figure 1, which shows the transformation of raw host observations into layer-specific scores and aggregated system-level risk metrics.

The Quantum-Adjusted Risk Scoring model was originally introduced in [35] as a multi-factor extension of Mosca's inequality formulated along three dimensions: timeline, sensitivity, and exposure. The present work reuses the QARS name, the three-dimensional decomposition, the weighted linear aggregation, and the Mosca-based time-budget grounding without modification. Three elements are new relative to [35]: (i) the scoring function is applied separately to in-transit, in-use, and at-rest data, producing layer-resolved scores rather than a single system-level score; (ii) a quantum-vulnerability attenuation pair (α, β) is introduced to scale time-based and impact-based factors when the dominant protection is Grover-bounded symmetric encryption; and (iii) the system-level output

is reported as a dual metric (Q_{agg}, Q_{max}) that separates overall posture from immediate remediation priority.

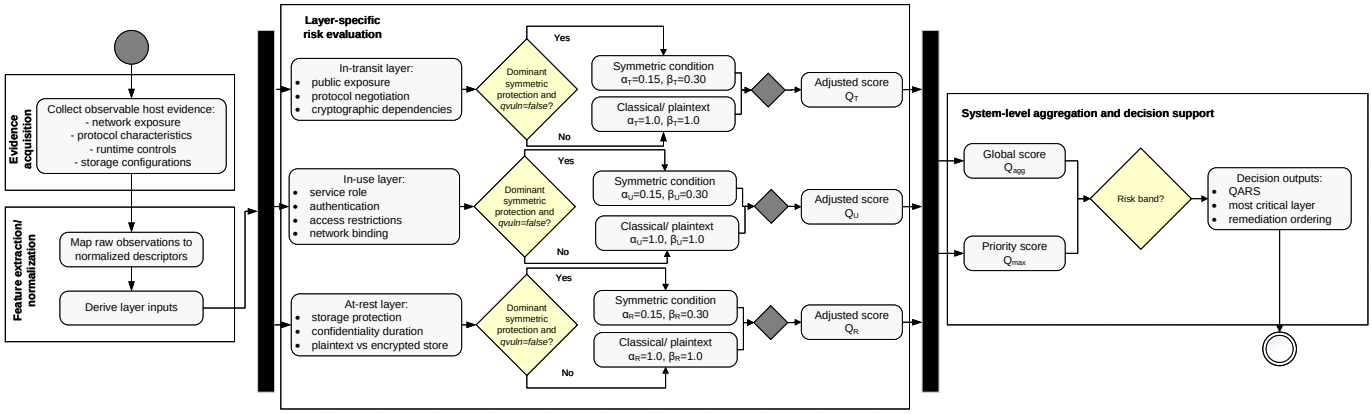

Figure 1. Workflow of the Quantum-Adjusted Risk Scoring model.

3.1. Model Overview

The proposed model follows a structured pipeline consisting of four main stages: (i) evidence acquisition, (ii) feature extraction, (iii) layer-specific scoring, and (iv) system-level aggregation. As shown in Figure 1, the process begins with the collection of observable host-level evidence, including network exposure, protocol characteristics, runtime controls, and storage configurations. This evidence is then transformed into normalized descriptors that serve as inputs to the scoring model.

Each data layer is evaluated independently to preserve its specific security properties. The resulting scores are subsequently aggregated to provide both a global system view and a prioritization-oriented perspective of post-quantum risk.

3.2. Canonical Scoring Model

The present work expands the original three-factor QARS formulation in [35] into a five-factor structure: the timeline dimension is split into a confidentiality duration factor r_1 and a migration effort factor r_2, and a compliance penalty C is added. This expanded form is then applied separately to each data layer, which is the principal extension introduced here. The QARS model builds upon a canonical five-factor risk formulation, where each layer score is computed as

$$Q_l = \text{clip}_{[0,1]}(w_1 r_{1,l} + w_2 r_{2,l} + w_3 S_l + w_4 E_l + w_5 C), \tag{1}$$

where $r_{1,l}$ represents the confidentiality duration factor, $r_{2,l}$ the migration effort factor, S_l the sensitivity, E_l the exposure, and C a compliance-related penalty term. The clipping operator ensures that the score remains within the interval $[0, 1]$.

The temporal factors are defined as

$$r_{1,l} = \min\left(\frac{T_l^{conf}}{T_Q}, 2\right), \quad r_{2,l} = \min\left(\frac{T_l^{mig}}{T_B}, 2\right), \tag{2}$$

where T_Q denotes the assumed quantum threat horizon and T_B represents the migration buffer period.

The weight vector $(w_1, w_2, w_3, w_4, w_5)$ is selected according to a critical-infrastructure profile and reflects the relative importance of time-based and impact-based risk components.

Following the critical-infrastructure profile implemented in the materials, the weight vector is

$$(w_1, w_2, w_3, w_4, w_5) = (0.25, 0.20, 0.20, 0.20, 0.15). \tag{3}$$

The planning constants are fixed at $T_Q = 12$ years, $T_B = 3$ years, and compliance penalty $C = 0.20$.

3.3. Quantum-Vulnerability Attenuation

The attenuation mechanism described in this subsection is new in the present work and has no counterpart in the original QARS formulation of [35], which does not distinguish between Shor-vulnerable and Grover-bounded protection. A key extension introduced in this work is the incorporation of a quantum-vulnerability attenuation mechanism. The model distinguishes between classical (quantum-vulnerable) and symmetric (non-Shor-vulnerable) protection mechanisms. The inventory auto-evaluator attenuates time-based factors by α and sensitivity/exposure factors by β whenever the dominant mechanism is marked as `qvuln = false`, with nominal values $\alpha = 0.15$ and $\beta = 0.30$.

The adjusted scoring function is defined as

$$\widetilde{Q}_l = \mathrm{clip}_{[0,1]}\left(w_1\alpha_l r_{1,l} + w_2\alpha_l r_{2,l} + w_3\beta_l S_l^{\mathrm{obs}} + w_4\beta_l E_l^{\mathrm{obs}} + w_5 C\right) \tag{4}$$

with

$$(\alpha_l, \beta_l) = \begin{cases} (1.0, 1.0), & \text{classical or plaintext condition,} \\ (0.15, 0.30), & \text{dominant symmetric protection condition.} \end{cases} \tag{5}$$

This adjustment ensures that layers protected by strong symmetric encryption (e.g., AES-based storage) are not assigned the same post-quantum urgency as layers relying on quantum-vulnerable public-key mechanisms. Such adjustment is important for the at-rest layer. A plaintext store remains fully exposed, whereas an AES-XTS-based LUKS2 layer should not inherit the same post-quantum urgency as an RSA- or ECDH-dependent transport channel.

Cryptographic Rationale for α and β

The choice of α = 0.15 and β = 0.30 is grounded in the asymmetric impact of quantum algorithms on symmetric versus asymmetric cryptography. Shor's algorithm [1] reduces the hardness of factoring and discrete logarithms to polynomial time, which renders RSA, ECDH, ECDSA, and finite-field DH cryptographically broken once a sufficiently large CRQC exists. Grover's algorithm [2], by contrast, provides only a quadratic speed-up against unstructured key search: a k-bit symmetric key offers an effective post-quantum work factor of $2^{k/2}$ rather than 2^k. For LUKS2 with the default `aes-xts-plain64` cipher and 512-bit key material (interpreted in XTS as 2×256-bit halves) [7,11], the per-block AES-256 key retains an effective post-quantum security level of approximately 128 bits, which remains comfortably above the 112-bit security floor recommended by NIST SP 800-131A [11] and the CNSA 2.0 advisory [10]. Even AES-128 in XTS mode yields a Grover-bounded floor of approximately 64 bits per block, which is below modern recommendations but still does not collapse to a zero work factor; Grover's algorithm is moreover difficult to parallelize linearly because the quadratic speed-up degrades to $\sqrt{T}$ across T machines, sharply increasing wall-clock cost.

The two coefficients α and β encode this asymmetry through different components of the canonical formula:

- $\alpha = 0.15$ scales the time-based factors $r_{1,l}$ and $r_{2,l}$. These factors capture the quantum-relevant time horizon (confidentiality duration vs. T_Q) and migration ur-

gency (effort vs. T_B). Because Grover's algorithm does not collapse symmetric security to a polynomial bound, the time-related migration pressure on appropriately keyed AES-XTS volumes is dominated by hash and KDF replacement rather than by an immediate decryption threat. Setting α to roughly one-sixth of the classical coefficient reflects the residual time-based pressure attributable to long-tail effects such as key-derivation algorithm transition (Argon2id parameter migration) and policy-driven recryption windows.

- $\beta = 0.30$ scales the impact-based factors S_l^{obs} and E_l^{obs}. The larger residual reflects the fact that symmetric encryption does not eliminate sensitivity- or exposure-driven risk: side-channel attacks, key-management failures, and key disclosure through cross-layer compromise (Section 3.11) remain pertinent even when the cryptographic primitive is Grover-bounded. Setting β to roughly one-third of the classical coefficient captures this irreducible residual while still distinguishing protected from unprotected configurations.

The $\alpha < \beta$ ordering is therefore a direct consequence of the different threat surfaces that the two factor families capture: time-based factors lose more of their post-quantum salience under symmetric protection than impact-based factors. The robustness of the resulting layer ranking under perturbations of these coefficients is examined empirically in the sensitivity analysis of Section 6.5.

The same coefficients are applied uniformly to dominant symmetric configurations within the evaluated profile. They are not intended to capture finer distinctions such as AES-128 vs. AES-256 keying or the choice between XTS and GCM modes; such distinctions are deliberately deferred to a future extended scoring model and are listed as a limitation in Section 7.6.

3.4. Evidence-to-Feature Mapping

The proposed model separates evidence extraction from coefficient calibration. Host-level observations determine the binary and ratio descriptors used in the model, including public exposure, legacy protocol presence, weak-cipher ratios, authentication state, access control state, proxy mediation, and disk encryption status. These observed descriptors are then mapped into normalized sensitivity, exposure, and migration effort variables through bounded affine rules. The coefficients in these rules are not claimed to be empirically fitted constants; they are transparent semi-quantitative calibration parameters whose provenance is documented in Section 3.4.2 and whose robustness is evaluated in Section 6.5. The mapping process, illustrated in Figure 1, converts raw observations into normalized layer-specific variables.

For the in-transit layer, factors such as legacy protocol presence, weak cipher ratios, and public exposure are used to derive sensitivity and exposure scores. For the in-use layer, runtime characteristics including service role, authentication, access control, and network binding are considered. For the at-rest layer, storage protection mechanisms and confidentiality requirements define the primary risk factors. This evidence-based transformation ensures auditability and reproducibility of the scoring process while maintaining alignment with real-world infrastructure conditions.

3.4.1. Generalized Observation Taxonomy

To support reproducible mapping beyond the Redis-on-Host 2 example used in this study, Table 2 presents a generalized taxonomy of observation classes for S^{obs} and E^{obs} across the three data layers. Each observation class is associated with a canonical evidence source (network scan, service configuration, host inventory) and a numerical contribution range. The taxonomy is aligned with the cryptographic asset classes defined in OWASP

CycloneDX CBOM [27], the cryptographic discovery practices in NIST SP 1800-38B [7], and the asset categorization in CARAF [25].

Table 2. Generalized observation taxonomy for evidence-to-feature mapping. Each class lists a canonical evidence source and the contribution range to S^{obs} or E^{obs}. The Redis/Host 2 case in this paper is one instantiation.

Layer	Observation Class	Canonical Evidence Source	Contributes to	Range
In-Transit	Public surface ratio	Port scan/external reachability test	E_T^{obs}	$[0,1]$
In-Transit	Legacy protocol indicator (TLS $\leq$ 1.1, SSHv1)	TLS handshake fingerprint/ banner grab	S_T^{obs}	$\{0,1\}$
In-Transit	Weak-cipher ratio (CBC, RC4, 3DES, MD5)	Cipher-suite enumeration	$S_T^{\text{obs}}, E_T^{\text{obs}}$	$[0,1]$
In-Transit	Public-key algorithm class	Certificate/key-exchange parameter inspection	T_T^{mig}	$\{0,1\}$
In-Use	Service criticality role	Service inventory/configuration	S_U^{obs}	$\{0,1\}$
In-Use	Authentication state	Configuration audit (`requirepass`, ACL, mTLS)	E_U^{obs}	$\{0,1\}$
In-Use	Access control state	Service configuration	E_U^{obs}	$\{0,1\}$
In-Use	Bind exposure (loopback vs. public)	`ss`/`lsof` or equivalent	E_U^{obs}	$\{0,1\}$
In-Use	Persistence enabled (write durability)	Configuration audit	S_U^{obs}	$\{0,1\}$
In-Use	Transport wrapping presence (stunnel, mTLS sidecar)	Process/network inventory	Modifier (see Classification of Redis-Related Evidence)	$\{0,1\}$
At-Rest	Storage encryption state	`cryptsetup`/`lsblk`/ FDE inventory	E_R^{obs} and (α, β) flag	$\{0,1\}$
At-Rest	Cipher and key length	`cryptsetup luksDump`/ FDE metadata	(α, β) flag	categorical
At-Rest	Confidentiality duration class	Data-classification policy/regulatory mapping	T_R^{conf}	years
At-Rest	Key-derivation function class	FDE metadata (Argon2id, PBKDF2 iterations)	Modifier on α	categorical

The aggregated forms of S_l^{obs} and E_l^{obs} are bounded affine combinations of the indicator and ratio variables in Table 2, clipped at unity. The Redis/Host 2 instantiation in Section 4 is therefore one realization of this general taxonomy and not a special case.

3.4.2. Coefficient Provenance and Calibration Status

The coefficients used in Equations (9)–(15) are semi-quantitative calibration parameters rather than empirically fitted constants. This treatment follows the general risk-modeling principle in NIST SP 800-30 Rev. 1: a risk assessment methodology should make explicit the risk factors, value scales, and algorithms used to combine those factors, while allowing organization-specific models to define their own combination rules. The coefficients in the present model therefore serve as a transparent and reproducible mapping from observed host evidence to normalized QARS inputs, not as universal constants.

The mapping follows the following design constraints. All coefficients are non-negative, so that additional adverse evidence cannot reduce the corresponding sensitivity, exposure, or migration effort score. Sensitivity and exposure outputs are clipped to [0,1], preventing additive evidence from producing values outside the model domain. The coefficients encode ordinal priority among evidence classes: public reachability, obsolete protocol support, weak cryptographic negotiation, missing authentication, missing access control, and absence of mediation are treated as stronger exposure drivers than purely contextual indicators. The constants represent that baseline layer risk by a transport surface, runtime

service, or storage path retains some residual sensitivity or exposure even when no single adverse indicator is present.

For the in-transit layer, the sensitivity rule assigns the largest contribution to the baseline transport condition, followed by legacy protocol support, weak-cipher prevalence, and public surface ratio. The corresponding exposure rule assigns the largest contribution to the baseline externally reachable communication surface, followed by public reachability, weak-cipher prevalence, and public SSH exposure. This reflects the fact that exposure is driven primarily by reachable attack surface, whereas sensitivity is driven by the protected data and cryptographic material carried by the service.

For the in-use layer, the Redis master role and append-only persistence increase sensitivity because they indicate active operational state and durable writes. Authentication and ACL presence contribute only small positive terms to sensitivity because they indicate that the service is protecting controlled data; their protective effect is modeled separately in Equation (14), where missing authentication and missing ACLs increase exposure. Public binding is assigned the largest exposure coefficient because direct non-loopback reachability is the strongest runtime exposure driver, while proxy mediation reduces exposure by ensuring that Redis is not reached directly.

The migration effort formulas in Equations (11) and (15) are low-resolution planning heuristics. Each active complexity condition adds up to 0.5 years: legacy protocol support, weak-cipher pressure, and proxy-mediated service wrapping increase the migration effort because they require compatibility testing, staged rollout, and fallback management. These values are therefore treated as calibration defaults and are stress-tested in Section 6.5.

3.5. Layer-Specific Risk Evaluation

Each data layer is evaluated independently to capture its distinct contribution to post-quantum risk:

- **In-transit:** Characterized by public exposure, protocol negotiation properties, and cryptographic dependencies.
- **In-use:** Determined by runtime controls, service roles, and access restrictions.
- **At-rest:** Driven by confidentiality duration and the strength of storage protection mechanisms.

This separation is essential because quantum vulnerability manifests differently across data states, particularly in the contrast between public-key-dependent communication channels and symmetrically protected storage systems.

3.6. System-Level Aggregation

To support both strategic and operational decision-making, the model defines two complementary system-level metrics. For whole-environment assessment, the aggregate is

$$Q_{agg} = 0.40\widetilde{Q}_T + 0.20\widetilde{Q}_U + 0.40\widetilde{Q}_R. \tag{6}$$

For near-term remediation ordering, the environment also reports

$$Q_{max} = \max\left(\widetilde{Q}_T, \widetilde{Q}_U, \widetilde{Q}_R\right). \tag{7}$$

The weighted aggregate score Q_{agg} provides an overall assessment of infrastructure risk, while the maximum-layer score Q_{max} identifies the most critical layer requiring immediate remediation.

Risk bands are unchanged: low for $Q < 0.35$, moderate for $0.35 \leq Q < 0.65$, and high for $Q \geq 0.65$.

3.7. Model Workflow Integration

The complete workflow of the proposed model integrates all stages from evidence acquisition to final risk evaluation. The flowchart highlights the sequential transformation of input data into actionable risk metrics and emphasizes the modular structure of the model.

This design ensures that the QARS framework remains adaptable to different infrastructure environments while preserving consistency in risk evaluation across heterogeneous systems.

3.8. Differentiation of QARS from Prior Quantum-Risk Frameworks

The QARS formulation in Equation (1) extends a canonical five-factor formulation that has its roots in two complementary lines of work: the Mosca timeline argument [28], which establishes the relationship $X + Y > Z$ between the data-confidentiality lifetime X, the migration time Y, and the time Z to a CRQC; and the CARAF framework [25], which decomposes crypto-agility risk into a 5-step pipeline (threat vector, asset inventory, asset value, mitigation, roadmap). Both formulations contribute the underlying notion of a time-driven scoring component normalized against a quantum-threat horizon, but neither produces a per-layer numerical score that distinguishes data states.

Three components of the present work are reused from these prior formulations:

1. The temporal factors $r_{1,l}$ and $r_{2,l}$ as ratios of T^{conf} and T^{mig} against T_Q and T_B, respectively, are the direct numerical analogue of the Mosca timeline inequality [28];
2. The linear weighted aggregation of confidentiality, migration, sensitivity, exposure, and compliance is structurally consistent with the asset-value step of CARAF [25];
3. The use of cryptographic asset inventories as input is consistent with NIST SP 1800-38B [7] and OWASP CycloneDX CBOM [27].

Three components are new in this work relative to the canonical and prior-art formulations:

1. The layer-specific evaluation in which the same scoring function is applied separately to in-transit, in-use, and at-rest data, producing layer-resolved scores that prior frameworks aggregate or omit;
2. The quantum-vulnerability attenuation mechanism of Equations (4) and (5), in which the (α, β) pair encodes the cryptographic distinction between Shor-vulnerable and Grover-bounded primitives;
3. The dual reporting of Q_{agg} for posture and $Q_{\max}$ for prioritization, which is not produced by CARAF, ETSI TR 104 016, or the NIST migration practice guides.

Operationally, the model differs from prior frameworks in two respects: the inputs are derived from host-level evidence rather than from organizational self-assessment questionnaires, and the outputs are numerical layer scores rather than ordinal risk categories. A side-by-side comparison against CARAF [25], ETSI TR 104 016 [8], NIST SP 1800-38 [7], the Mosca timeline argument [28], FS-ISAC PQC migration recommendations as summarized in industry reports [13], and CycloneDX CBOM [27] is presented as part of the experimental evaluation in Section 6.6.

3.9. Parameter Calibration and Justification

The parameters below combine inherited elements from [35] (sector-specific weight calibration as a principle, as well as the Mosca-based horizon T_q) with new elements introduced here (T_B, C, the attenuation pair (α, β), and the aggregation weights for Q_{agg}). The five-factor weight vector w is calibrated independently from the original three-factor vector (w_T, w_S, w_E). The model uses six families of fixed parameters: the quantum threat horizon T_Q, the migration buffer T_B, the compliance penalty C, the five-factor weight vector

w, the quantum-vulnerability attenuation pair (α, β), and the aggregation weights for Q_{agg}. Each is assigned a defensible default value drawn from public guidance, with the explicit assumption that operators may re-calibrate to local conditions; the sensitivity of the final scores to these defaults is examined in Section 6.5.

Quantum threat horizon T_Q.

The default $T_Q = 12$ years follows the most-cited median estimate of the Global Risk Institute Quantum Threat Timeline series [36], which reports a median expert estimate of approximately 11–14 years for the appearance of a CRQC capable of breaking RSA-2048. This value is also consistent with the planning horizons referenced in NIST SP 1800-38A [7], the UK NCSC migration timelines [12], and the Mosca $X + Y > Z$ argument [28]. We treat T_Q as an institutional planning constant rather than as an empirical estimate, and we accept its inherent uncertainty as part of the modeling assumption.

Migration buffer T_B.

The default $T_B = 3$ years reflects the typical large-organization cryptographic migration cycle reported by NIST SP 1800-38A [7], ETSI TR 103 619 [9], and post-mortems of prior cryptographic transitions such as SHA-1 deprecation [13,25]. Industrial migration retrospectives identify three years as the median from formal program kickoff to fleet-wide rotation in regulated environments.

Compliance penalty C.

The default $C = 0.20$ reflects an additive risk floor for organizations subject to mandatory cybersecurity regulation in the EU. It is calibrated to the proportional weight of cryptographic compliance obligations under NIS2 [30], the EU Cyber Resilience Act, and sector-specific banking and utility directives. The value is conservative: an organization under no regulatory regime would set C closer to 0, while an organization in a heavily regulated critical-infrastructure context could justify C as high as 0.30.

Five-factor weight vector w.

The vector $(0.25, 0.20, 0.20, 0.20, 0.15)$ implements a critical-infrastructure profile in which time-based factors slightly outweigh impact-based factors, matching the priorities expressed in NIS2 [30] and the NSA CNSA 2.0 advisory [10]. Two alternative profiles are reported in the sensitivity analysis: a transit-prioritized profile and an impact-prioritized profile.

Quantum-vulnerability attenuation $(\alpha, \beta) = (0.15, 0.30)$.

The cryptographic justification is provided in subsection "Cryptographic Rationale for α and β". The values are not derived from experiment but from the qualitative ordering Shor $\gg$ Grover applied to the time-related and impact-related factor families. The robustness of layer ranking under perturbations $\alpha \in [0.075, 0.225]$ and $\beta \in [0.15, 0.45]$ is examined in Section 6.5.

Aggregation weights $(0.40, 0.20, 0.40)$.

Transport and storage receive equal weights, reflecting their equally severe but qualitatively different post-quantum exposure: transport is dominated by Shor-vulnerable key exchange, while at-rest storage is dominated by long confidentiality horizons. The in-use layer receives a smaller weight (0.20) because runtime exposure is typically controllable through operational hardening rather than through cryptographic transition. Sensitivity to alternative aggregation profiles is reported in Section 6.5.

Risk thresholds 0.35 and 0.65.

The thresholds partition the unit interval into three approximately equal-width bands, in line with the three-category risk reporting used in CARAF [25], OWASP risk-rating practice, and NIS2 supervisory guidance [30]. The 0.65 boundary corresponds to a scoring profile in which at least one of the two highest-weight factors is at or near its maximum.

3.10. Input Provenance: Observed, Assumed, and Calibrated Parameters

Not every input to QARS is directly observed; some are calibrated and others are assumed. To make this distinction explicit and auditable, Table 3 classifies every parameter into one of three categories:

- **Observed**: The value is computed from a host scan, configuration file, or service inventory, and an evidence trail can be reconstructed from raw outputs.
- **Calibrated**: The value is fixed by a published reference (NIST, ETSI, NSA, IETF, or peer-reviewed) and inherits the uncertainty of that reference; it is intended to be re-tuned by the operator.
- **Assumed**: The value is set by the authors as a modeling choice without a single dominant external reference.

Table 3. Provenance of every QARS parameter, classified as Observed, Calibrated, or Assumed.

Parameter	Class	Source/Justification	Reference
$E_T^{\text{obs}}, S_T^{\text{obs}}$	Derived: observed descriptors + calibrated mapping	Port scan, TLS fingerprint, weak-cipher enumeration, SSH banner	Host 1 ports sheet
$E_U^{\text{obs}}, S_U^{\text{obs}}$	Derived: observed descriptors + calibrated mapping	Redis configuration, bind address, persistence flag, ACL state	Host 2 ports sheet
E_R^{obs} (plaintext)	Observed	`lsblk/cryptsetup` reports no FDE	Host 1 disk sheet
E_R^{obs} (protected)	Observed	LUKS2 detected, AES-XTS-plain64, 512-bit material	Host 2 disk sheet
T_T^{conf}	Calibrated	Standard TLS-transport heuristic (5 years)	[7,8]
T_U^{conf}	Assumed	Operational state typically short-lived (2 years)	this work
T_R^{conf}	Calibrated	Long-term archival horizon (15 years)	[8,25]
T_T^{mig}	Calibrated heuristic	Heuristic on legacy and weak-cipher pressure	this work, anchored in [7]
$T_U^{\text{mig}}, T_R^{\text{mig}}$	Calibrated heuristic	Disk-encryption rest heuristic (3 years)	[7,13]
$T_Q = 12$ years	Calibrated	Median GRI quantum-threat estimate	[28,36]
$T_B = 3$ years	Calibrated	Standard large-org migration cycle	[7,9]
$C = 0.20$	Assumed	Critical-infrastructure compliance floor	[30] (range guidance)
$w = (0.25, 0.20, 0.20, 0.20, 0.15)$	Calibrated	Critical-infrastructure profile	[25,30]
$(\alpha, \beta) = (0.15, 0.30)$	Calibrated	Grover-bounded residual security argument	[1,2,11]
Aggregation weights	Assumed	Transit/at-rest equal, in-use lower	this work
Thresholds 0.35, 0.65	Calibrated	Three-band risk partition	[25]

This classification responds directly to the criticism that some quantities described as "derived from host evidence" are in fact heuristics. The classification informs the sensitivity analysis (Section 6.5) and supports operator re-calibration.

3.11. Cross-Layer Interdependencies

The QARS scoring function evaluates each layer independently, which is consistent with the divide-and-conquer migration philosophy of ETSI TR 104 016 [8] and is well suited to producing actionable per-layer remediation priorities. However, layer independence is a modeling simplification: in practice, a compromise in one layer can cascade into another, and a comprehensive risk picture must acknowledge these couplings.

Three classes of cross-layer dependency are particularly relevant to post-quantum risk:

- **In-use → at-rest leakage.** A memory disclosure or arbitrary-read vulnerability in a service running on top of an encrypted volume effectively bypasses LUKS2 protection at the application boundary. The at-rest layer is then only as strong as the in-use layer protecting the unsealed key material. This is the dominant cross-layer channel for runtime services such as Redis, where keys exist in cleartext in the process memory while the service is running.
- **In-transit → in-use compromise.** A successful TLS-layer attack (e.g., a downgrade to a Shor-vulnerable suite) can yield session keys whose loss exposes the runtime service. Transport-layer authentication failures degrade in-use isolation regardless of the in-use layer's local controls.
- **At-rest → in-transit/in-use unsealing.** TLS server private keys, SSH host keys, and OAuth signing material are typically stored at rest. A compromise of the at-rest key vault therefore exposes both transport and runtime boundaries, even if those layers report low immediate risk.

These couplings imply that the layer scores are first-order independent but second-order interdependent. Operationally, this argues that the dominant-layer remediation priority indicated by $Q_{\max}$ should be implemented in conjunction with hardening of the layers that gate access to its key material. Formally extending QARS to a coupled-layer model would require introducing a per-pair coupling matrix; this is left as future work.

4. Illustrative Case Study: Source Data and Extracted Evidence

The experimental evaluation of the proposed QARS model is based on structured host-level evidence collected from a controlled infrastructure dataset. The dataset consists of four workbook sheets, namely `Host_1_ports_v8`, `Host_1_disk_v8`, `Host_2_ports_v8`, and `Host_2_disk_v8`. Table 4 summarizes the evidence used in the update.

The two hosts are not treated as independent systems but as complementary evidence sources representing different security conditions within a single infrastructure context. Host 1 serves as the baseline reference for externally exposed communication surfaces and unprotected storage conditions, while Host 2 represents a more controlled operational configuration with restricted runtime exposure and encrypted storage. This design enables a comparative evaluation of post-quantum risk under different protection scenarios while maintaining consistency across unaffected layers.

The data extraction process follows an evidence-driven approach in which raw workbook observations are transformed into structured descriptors suitable for model input. Instead of directly using raw values, the extraction procedure identifies security-relevant characteristics such as publicly reachable services, protocol versions, cryptographic configurations, authentication mechanisms, access control settings, and storage protection mechanisms. These characteristics are then normalized into intermediate descriptors that reflect the exposure, sensitivity, and operational constraints of each data layer.

For the in-transit layer, the extracted evidence captures externally accessible services, protocol negotiation properties, and cryptographic weaknesses, including legacy protocol usage and weak cipher configurations. For the in-use layer, the extraction focuses on runtime properties such as service role, binding configuration, authentication enforcement, access control mechanisms, and the presence of proxy-mediated communication. For the at-rest layer, the process identifies storage protection mechanisms, including the presence or absence of disk encryption, encryption schemes, and key-derivation configurations.

A key feature of the extraction process is its reliance on observable infrastructure evidence rather than synthetic assumptions. This ensures that the derived model inputs remain auditable and reproducible, while accurately reflecting real-world system configurations. The extracted descriptors are subsequently used as inputs to the QARS model, where they are mapped to layer-specific parameters for confidentiality duration, migration effort, sensitivity, and exposure.

The separation of evidence extraction from scoring enables a modular evaluation workflow. First, infrastructure observations are collected and transformed into standardized descriptors. Second, these descriptors are mapped to quantitative parameters. The resulting parameters are used within the scoring model to evaluate post-quantum risk across data layers. This structured approach ensures that the experimental evaluation remains consistent, transparent, and adaptable to different infrastructure environments.

Table 4. Host evidence extracted from the workbook and used for the updated experiment.

Layer	Source	Key Observations
In-Transit	Host 1, ports sheet	Open services on ports 22, 25, 389, 443, and 636. Four TLS-capable services (25, 389, 443, 636). Weak ciphers on 3/4 TLS services. Legacy TLS 1.0 observed on SMTP. RSA-2048 certificates on four TLS services. Public SSH exposure present.
In-Use	Host 2, ports sheet	Redis 7.0.15 in master role. Bound to `127.0.0.1`. Protected mode enabled. Authentication required and used. `requirepass` configured. ACL enabled. Append-only persistence enabled. Approximate logical key count: 500.
In-use transport wrapping	Host 2, disk summary	Stunnel present with external accept endpoints `0.0.0.0:443` and `0.0.0.0:6380`, forwarding to `127.0.0.1:8080` and `127.0.0.1:6379`.
At-Rest control	Host 1, disk summary	No LUKS device detected. Redis-at-rest state unknown because no Redis directory mapping is present in the sheet. This condition is treated as the plaintext storage control.
At-Rest protected	Host 2, disk summary	LUKS2 present on `/dev/sda3`, mapped through `dm_crypt-0`. Cipher reported as `aes-xts-plain64` with 512-bit key material and Argon2id key derivation. Redis directory `/var/lib/redis` resides on the encrypted root mapper; the sheet explicitly marks effective at-rest protection as true.

4.1. Evidence Mapping from Workbook Observations

This paper does not insert raw workbook values directly into scoring. Instead, the workbook is transformed into observation scores that remain auditable. The mapping rules for each layer are presented in turn below.

4.2. Transit Layer Mapping

Let L_T be a legacy-TLS indicator, W_T the fraction of TLS services with weak ciphers, P_T the normalized public surface ratio, and S_{ssh} a public-SSH exposure indicator. Using Host 1,

$$L_T = 1, \quad W_T = \frac{3}{4} = 0.75, \quad P_T = \min\left(\frac{5}{5}, 1\right) = 1.000, \quad S_{ssh} = 1. \tag{8}$$

The observed sensitivity and exposure are then defined as

$$S_T^{\text{obs}} = \min\left(0.45 + 0.20L_T + 0.15W_T + 0.10P_T,\ 1\right) = 0.862, \tag{9}$$

$$E_T^{\text{obs}} = \min\left(0.55 + 0.20P_T + 0.15W_T + 0.10S_{ssh}, 1\right) = 0.963. \tag{10}$$

The confidentiality horizon follows the TLS transport heuristic, $T_T^{conf} = 5$, and migration time is elevated by legacy and weak-cipher pressure, as follows:

$$T_T^{mig} = 1.5 + 0.5L_T + 0.5W_T = 2.375. \tag{11}$$

4.3. In-Use Layer Mapping

For Redis on Host 2, let M_{role} denote master role criticality, P_{aof} persistent write enablement, A_{auth} successful authentication hardening, A_{acl} ACL enablement, B_{pub} non-loopback bind exposure, and Z_{proxy} transport wrapping presence. The workbook gives

$$M_{role} = 1, \quad P_{aof} = 1, \quad A_{auth} = 1, \quad A_{acl} = 1, \quad B_{pub} = 0, \quad Z_{proxy} = 1. \tag{12}$$

Thus

$$S_U^{\text{obs}} = \min\left(0.45 + 0.10M_{role} + 0.10P_{aof} + 0.05A_{auth} + 0.05A_{acl},\ 1\right) = 0.750, \tag{13}$$

$$E_U^{\text{obs}} = \min\left(0.20 + 0.35B_{pub} + 0.20(1 - A_{auth}) + 0.15(1 - A_{acl}) + 0.10(1 - Z_{proxy}),\ 1\right) = 0.200. \tag{14}$$

Because Redis holds a short-lived operational state rather than archival records, the experiment fixes $T_U^{conf} = 2$ years and

$$T_U^{mig} = 1.5 + 0.5Z_{proxy} = 2.000. \tag{15}$$

Classification of Redis-Related Evidence

The Redis service on Host 2 generates evidence that touches three layers simultaneously: (i) Redis itself is a runtime service with in-use security characteristics, (ii) stunnel forwarding through external endpoints `0.0.0.0:443` and `0.0.0.0:6380` introduces a transport surface, and (iii) the append-only persistence of Redis state writes to the encrypted volume creates an at-rest channel. To prevent double counting, this work classifies Redis-related evidence as follows:

- The **Redis configuration itself** (master role, loopback bind, authentication, ACL, persistence flag) contributes only to the in-use layer through S_U^{obs} and E_U^{obs}. This is the only contribution that enters $\widetilde{Q}_U$.
- The **stunnel external endpoints** are recorded for context as runtime mediation, captured exclusively through the Z_{proxy} indicator that reduces E_U^{obs} in Equation (14). They do not enter the in-transit layer score because the in-transit layer is computed from Host 1 and is intentionally held constant across scenarios as the experimental control.

https://doi.org/10.3390/electronics15122546

- The **append-only persistence directory** `/var/lib/redis` is treated as part of the at-rest layer, but only insofar as the encrypted-volume status determines the at-rest condition (plaintext vs. LUKS2). It does not enter the in-use score.

This deliberate one-evidence–one-layer routing eliminates double counting and, in particular, prevents stunnel evidence from contributing simultaneously to in-transit and in-use scores. The Z_{proxy} indicator is best understood as a runtime-mediation modifier that lowers in-use exposure because external traffic is mediated through a controlled gateway rather than reaching Redis directly.

4.4. At-Rest Layer Mapping

The at-rest layer uses one sensitivity value for both conditions, $S_R^{\text{obs}} = 0.85$, and one planning pair, $T_R^{conf} = 15$ years and $T_R^{mig} = 3$ years, following the disk-encryption rest heuristic. Exposure differs by observed condition as follows:

$$E_{R,\text{plain}}^{\text{obs}} = 1.00, \qquad E_{R,\text{prot}}^{\text{obs}} = 0.350. \tag{16}$$

The plaintext control uses $(\alpha, \beta) = (1, 1)$. The LUKS2 condition uses $(\alpha, \beta) = (0.15, 0.30)$ because the dominant control is symmetric AES-XTS-based storage encryption.

4.5. Measurement Uncertainty and Robustness of Observations

Indicator variables such as the weak-cipher ratio W_T, the public surface ratio P_T, the legacy-TLS indicator L_T, and the storage-protection state depend on the scanning method, scanner version, and time of collection. To avoid overclaiming the precision of the resulting scores, three categories of uncertainty are explicitly recognized:

- **Detection uncertainty.** Network scanners may miss services behind connection rate-limiting, port-knocking, or stateful firewalls. Conversely, banner mismatches may misclassify TLS versions when servers downgrade-fingerprint as 1.2 while accepting 1.0 at the protocol layer. We treat the indicator variables as having an uncertainty band of ± 1 binary count per indicator class within the evaluated profile.
- **Temporal drift.** A scan is a snapshot. Cipher suite and protocol-version state can change with operating system updates or load-balancer reconfigurations on the order of weeks. The scoring rules treat the most recent valid scan as ground truth, with the implicit understanding that re-scoring is required after any cryptographic configuration change.
- **Categorical mapping uncertainty.** The mapping from raw configuration values to indicator variables (e.g., assigning $L_T = 1$ if any TLS service exhibits any version ≤ 1.1) is a deliberate aggregation. Operators can refine this by replacing the binary L_T with a graded indicator without changing the structure of Equations (9)–(11).

Section 6.5 examines the propagation of these uncertainties into the layer scores by perturbing the observation indicators and recomputing $\widetilde{Q}_T$, $\widetilde{Q}_U$, and $\widetilde{Q}_R$. The qualitative finding–that the transport layer dominates after at-rest hardening–is found to be robust to one-step perturbations of any single indicator class.

5. Case Study Design

The experiment was designed as an evidence-based three-layer assessment of post-quantum exposure in a national infrastructure setting. A high-level conceptual infrastructure architecture shows three main layers: the user layer, the data transmission layer, and the infrastructure layer (Figure 2).

The user layer shows user and device interaction; in this layer, data is actively in use. The transmission layer shows the movement of information across public network

and secure communication channels–data is in transit. The infrastructure layer contains two additional layers. The first one is the service and data processing layer, where data is processed. This layer represents data in use. The second is the data storage layer; in this layer, data is maintained in databases and file and object storage systems, representing the at-rest state. Bidirectional arrows between storage components and the processing component show the constant movement of data–reading, writing, and storing information.

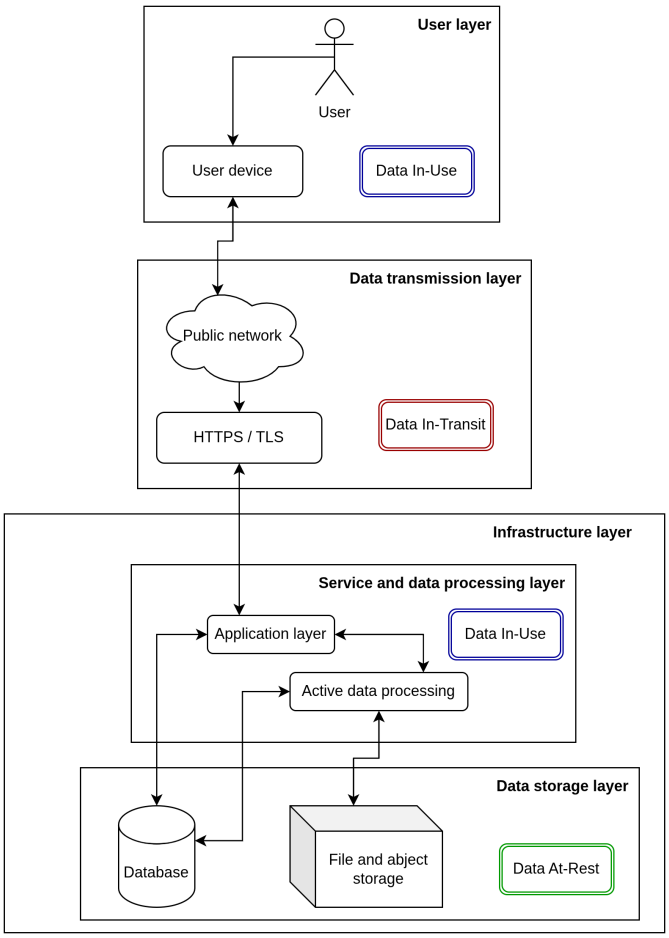

Figure 2. Architecture of the infrastructure for QARS application.

Instead of using synthetic assumptions, this study derived all its scoring inputs from host observations recorded in the supplied workbook. The environment was decomposed into in-transit, in-use, and at-rest layers because quantum-related exposure is not uniform across the data lifecycle. Communication channels are driven by public reachability, protocol age, and cryptographic negotiation properties. Active operational processing is influenced by service role and runtime controls. Stored data is driven by confidentiality duration and the strength of the protecting mechanism. Each layer was therefore evaluated separately and then recombined into a system-level result. This structure ensured that the experiment measured both total risk and the location of the dominant migration bottleneck.

The empirical basis of the experiment consisted of four workbook sheets describing two hosts. Host 1 was treated as the baseline source for public communication exposure and the unprotected storage condition. Host 2 was treated as the protected operational source for active-state processing and encrypted storage. The hosts were not interpreted as unrelated systems. They were used as complementary evidence sources in a controlled comparative design. This allowed this study to hold two layers constant while modifying only one layer between scenarios. The resulting experiment is best described as a within-

case comparative design with two scenarios: a plaintext-storage control condition and a protected-storage condition.

The evaluation reported in this paper is structured as an illustrative case study rather than a statistical validation. Its purpose is to demonstrate the evidence-to-score pipeline end-to-end on real host observations, to show that the layer-resolved scores respond meaningfully to a concrete remediation actions, and to expose the model's behavior under sensitivity perturbations of its calibration parameters. The case study is not designed to support claims about the distribution of QARS scores across a production fleet, nor the statistical generalizability of the layer ordering observed here. Those claims would require a sample of independently configured hosts, which is outside the scope of the present work and is discussed as a future work direction in Section 7.6. The robustness evidence reported below is therefore of two kinds: structural robustness against perturbations of the calibration parameters (Section 6.5) and qualitative consistency with the cryptographic reasoning of subsection "Cryptographic Rationale for α and β".

The in-transit layer was instantiated from Host 1 and represented externally reachable communication surfaces through which encrypted sessions could be intercepted, downgraded, or harvested for later decryption. The workbook indicated open services on ports 22, 25, 389, 443, and 636, corresponding to SSH, SMTP, LDAP, HTTPS, and LDAPS. Four services were considered TLS-capable, three of them exhibited weak-cipher observations, and one legacy TLS 1.0 condition was present. Public SSH exposure was also retained because it enlarged the cryptographic attack surface. This layer was intentionally modeled as a mixed legacy environment rather than an idealized modern deployment. The goal was to capture a realistic communication profile in which service continuity and backward compatibility could still coexist with incomplete cryptographic modernization.

5.1. Experimental Equipment and Environment

The evaluation was conducted on two Ubuntu Server 24.04.4 LTS virtual machines hosted on a single Proxmox VE 8.4.14 hypervisor, isolated from production traffic. Both VMs were provisioned with identical resources (2 vCPU, 8 GB RAM, 32 GB virtual disk) so that observed differences in risk scoring originated from guest security configuration rather than from hardware asymmetry. Host 1 served as the baseline (public communication surface, plaintext storage), and Host 2 as the protected condition (operational Redis processing, LUKS2-encrypted root). Evidence was collected with nmap 7.94, cryptsetup 2.7.0, redis-cli 7.0.15, and standard util-linux tools. Software on the targets included OpenSSH 9.6p1 and OpenSSL 3.0.13 on both hosts, as well as Redis 7.0.15, stunnel, and an LUKS2 volume with an AES-XTS-plain64 cipher and Argon2id key derivation on Host 2.

5.2. Plaintext Control Scenario

The transit layer was evaluated from Host 1 public transport evidence, the in-use layer from Host 2 Redis evidence, and the at-rest layer from Host 1 plaintext storage. The in-use layer was instantiated from Host 2 and modeled active processing through Redis. The workbook identified Redis 7.0.15 in the master role, bound to `127.0.0.1`, with protected mode enabled, authentication enabled, access control lists enabled, and append-only persistence enabled. The same host also used stunnel to forward external endpoints to local services. This detail prevented a false binary interpretation of Redis as either public or isolated. The service was modeled instead as locally constrained but operationally reachable through a mediation layer. This reflects a realistic pattern for critical environments in which sensitive runtime data is processed internally while business access is exposed through a controlled gateway.

The at-rest layer was the intervention variable of the experiment. It was evaluated under two conditions. In the control condition, storage protection was taken from Host 1, where no LUKS device was observed and the storage state was treated as plaintext for scoring purposes. In the protected condition, storage protection was taken from Host 2, where the workbook identified a LUKS2-backed encrypted root mapper protecting the Redis data path. The protection used AES-XTS-plain64 with 512-bit key material and Argon2id-based key derivation. Only this layer was changed between scenarios, while the transport and in-use layers were kept fixed. This control logic made it possible to measure the marginal effect of storage encryption on the total quantum-adjusted risk profile.

5.3. Protected Condition Scenario

The transit and in-use layers remain the same, but the at-rest layer is replaced by Host 2 LUKS2-protected storage. This design isolates the effect of storage protection without hiding the fact that transport risk remains unchanged. It is therefore appropriate for showing both absolute risk and marginal hardening benefit. The scoring stage followed the canonical structure and the three-layer logic, but it updated the calculation in two ways. First, confidentiality duration, migration time, sensitivity, and exposure were mapped from observed host evidence rather than assigned as abstract defaults. Second, the experiment introduced quantum-vulnerability attenuation for dominant symmetric protection states, following the supplied inventory logic. The critical-infrastructure weight set was retained, together with a twelve-year threat horizon and a three-year migration buffer. Classical and plaintext conditions were evaluated without attenuation. The protected at-rest condition, because it relied on symmetric storage encryption rather than a Shor-vulnerable public-key mechanism, was evaluated with reduced time-based and impact-based scaling. This prevented the model from overstating the post-quantum urgency of encrypted storage when compared with legacy public-key communication channels.

Execution proceeded in a fixed analytical sequence. Workbook observations were first normalized into measurable descriptors, such as public service count, weak-cipher ratio, legacy protocol presence, Redis role, local binding, authentication state, access control enablement, and encrypted volume presence. These descriptors were then converted into layer variables and entered into the updated QARS formula. Scores were computed separately for in-transit, in-use, and at-rest data, and then combined in two ways: a weighted aggregate score representing whole-system urgency and a maximum-layer score representing the most urgent remediation target. The design intentionally reported both measures because infrastructure decision-making requires a distinction between overall risk reduction and the persistence of a dominant unresolved layer. In practical terms, the experiment was structured to show whether improving storage protection materially lowers the aggregate score while still revealing whether legacy transport remains the main barrier to post-quantum readiness.

5.4. Reproducibility Artifacts

To support replication of all reported scores, the following artifacts are made available alongside this paper:

- An anonymized version of the four-sheet workbook (`Hosts_info_anon.xlsx`) containing the redacted Host 1 and Host 2 evidence used in this study, with hostnames, IP addresses, and certificate fingerprints removed and service banners truncated.
- A self-contained Python notebook (`qars_recompute.ipynb`) that loads the workbook, computes Equations (8)–(16) step by step, and reproduces every numerical entry in Tables 5 and 6 as well as Figure 3.

- The collection commands used to obtain raw evidence on each host, `nmap -sV -p- --script ssl-enum-ciphers`, `lsblk -f`, `cryptsetup luksDump`, `redis-cli CONFIG GET *`, and `ss -tlnp`, with sample output excerpts redacted to remove identifying information.
- A worked-example walkthrough recomputing $\widetilde{Q}_T$ from the raw workbook to the final score, suitable for use as a teaching example or as an audit trail.

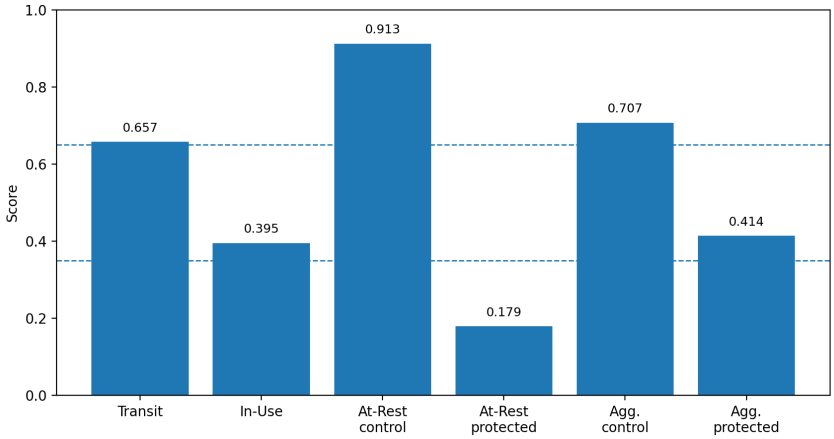

Figure 3. Updated layer and system scores derived from `Hosts_info.xlsx`. Dashed lines indicate the moderate threshold at 0.35 and the high threshold at 0.65.

Table 5. Derived layer parameters and intermediate factors.

Layer	T^{conf}	T^{mig}	S^{obs}	E^{obs}	α	β
Transit	5.000	2.375	0.862	0.963	1.000	1.000
In-Use	2.000	2.000	0.750	0.200	1.000	1.000
At-Rest Plaintext	15.000	3.000	0.850	1.000	1.000	1.000
At-Rest Protected	15.000	3.000	0.850	0.350	0.150	0.300

Table 6. System-level results for the plaintext control and the protected condition. All values rounded half-up to three decimals.

Scenario	$\widetilde{Q}_T$	$\widetilde{Q}_U$	$\widetilde{Q}_R$	Overall	Risk Band
Plaintext Control, weighted Q_{agg}	0.658	0.395	0.913	0.707	High
Protected Condition, weighted Q_{agg}	0.658	0.395	0.179	0.414	Moderate
Plaintext Control, max Q_{max}	0.658	0.395	0.913	0.913	High
Protected Condition, max Q_{max}	0.658	0.395	0.179	0.658	High

The artifacts are intended to be deposited in a public Git repository hosted on GitHub at the time of publication, in compliance with the host institution's data sovereignty and cybersecurity policies. Until that point, the artifacts are available from the corresponding author upon reasonable request, in line with the data availability statement at the end of this paper. The authors note that the original incompleteness rationale for restricted access has now been narrowed to identifying material only; the computational substance is fully replicable from the artifacts above.

6. Results

6.1. Derived Layer Parameters

The evaluation results are based on the parameter values derived from the evidence-mapping process. Table 5 summarizes the final inputs used for each data layer, including the confidentiality duration, migration time, sensitivity, exposure, and attenuation factors.

The derived values reflect distinct security characteristics across the three data states. The in-transit layer exhibits high exposure due to publicly accessible services and legacy protocol conditions, while the in-use layer shows reduced exposure as a result of enforced runtime controls. The at-rest layer demonstrates the largest variation between scenarios, depending on the presence or absence of storage encryption.

6.2. Layer-Level Results

The computed layer scores reveal a clear differentiation in post-quantum risk across data states. To support reproducibility, the intermediate temporal factors $r_{1,l}$ and $r_{2,l}$ are reported explicitly for each layer in Table 7, alongside the layer-level scores $\widetilde{Q}_l$. All numerical values are reported to three decimals; rounding is applied only at the display stage and is not propagated into subsequent computations, which use full double precision. In particular, the transport-layer score is $\widetilde{Q}_T = 0.65750$ before rounding; we display it as 0.658.

Table 7. Intermediate temporal factors and layer scores. All values rounded half-up to three decimals.

Layer	T^{conf}/T_Q	T^{mig}/T_B	r_1	r_2	Pre-Clip Sum	$\widetilde{Q}_l$
Transit	5/12	2.375/3	0.417	0.792	0.658	0.658
In-Use	2/12	2/3	0.167	0.667	0.395	0.395
At-Rest (plaintext)	15/12	3/3	1.250	1.000	0.913	0.913
At-Rest (LUKS2-protected)	15/12	3/3	1.250	1.000	0.179	0.179

In-transit layer. The layer score is obtained via direct substitution into Equation (4), as follows:

$$\widetilde{Q}_T = 0.25(0.417) + 0.20(0.792) + 0.20(0.862) + 0.20(0.963) + 0.15(0.20) = 0.658. \quad (17)$$

The transit result is high because the communication profile still contains several characteristics associated with quantum-vulnerable exchange and traffic capture risk. The scan evidence shows five active communication services on Host 1, including SSH on port 22 and TLS-bearing services on ports 25, 389, 443, and 636. Among the TLS-capable services, three present weak-cipher evidence, and the SMTP service includes a legacy TLS 1.0 condition. The certificate profile is RSA-2048-based, so the transport layer still depends on classical public-key mechanisms. Under the updated model, this combination pushes the transit score to 0.658, just above the high-risk boundary at 0.65. This matters because it shows that the communication layer is the limiting factor for post-quantum readiness.

In-use layer. The in-use result is more controlled and lands in the moderate band at 0.395, as follows:

$$\widetilde{Q}_U = 0.25(0.167) + 0.20(0.667) + 0.20(0.750) + 0.20(0.200) + 0.15(0.20) = 0.395. \quad (18)$$

This layer was derived from the Redis host and shows a mixed picture. Redis is configured on `127.0.0.1`, protected mode is enabled, authentication is required, access control lists are enabled, and append-only persistence is active. These settings reduce direct runtime exposure and explain why the score stays well below the transit result. At the same time, the service remains operationally reachable through stunnel bindings that accept external connections and forward them to local services; per subsection "Classification of Redis-Related Evidence", this presence enters the score only through the Z_{proxy} runtime-mediation modifier and not as a separate transport-layer contribution. The host also runs

Redis in the master role, so it remains part of an active processing path. For these reasons, the in-use layer is not low-risk, but it is clearly better controlled than the transport layer.

At-rest layer. The storage layer exhibits the most significant contrast between the evaluated scenarios. In the plaintext condition, the score reaches $Q_{R,plain} = 0.913$, representing the highest risk among all layers, as follows:

$$\widetilde{Q}_{R,\text{plain}} = 0.25(1.250) + 0.20(1.000) + 0.20(0.850) + 0.20(1.000) + 0.15(0.20) = 0.913. \quad (19)$$

This is primarily due to long confidentiality requirements combined with the absence of cryptographic protection. In contrast, the protected condition yields a substantially reduced score of $\widetilde{Q}_{R,prot} = 0.179$, as follows:

$$\widetilde{Q}_{R,\text{prot}} = 0.25(0.188) + 0.20(0.150) + 0.20(0.255) + 0.20(0.105) + 0.15(0.20) = 0.179, \quad (20)$$

where the bracketed quantities show the attenuated factors $\alpha r_1 = 0.15 \times 1.250 = 0.188$, $\alpha r_2 = 0.15 \times 1.000 = 0.150$, $\beta S^{\text{obs}} = 0.30 \times 0.850 = 0.255$, and $\beta E^{\text{obs}} = 0.30 \times 0.350 = 0.105$. The introduction of LUKS2-based AES-XTS-plain64 encryption with Argon2id key derivation, combined with the application of quantum-vulnerability attenuation, reduces both time-based and exposure-related risk contributions by approximately 80%. This demonstrates the effectiveness of symmetric encryption in mitigating post-quantum risk for stored data when the dominant control is Grover-bounded rather than Shor-vulnerable.

6.3. System-Level Results

The system-level evaluation combines layer scores using both aggregate and maximum-layer perspectives. Table 6 reports both the weighted aggregate score and the max-layer score.

The results show a clear shift in the dominant source of quantum-adjusted risk after storage protection is introduced. In the plaintext control scenario, the weighted aggregate score is $Q_{agg} = 0.707$, placing the system in the high-risk category. After introducing storage protection, the aggregate score decreases to $Q_{agg} = 0.414$, corresponding to a moderate-risk classification. This represents an absolute reduction of approximately 0.293, or 41.5%, indicating a substantial improvement in overall security posture.

However, the maximum-layer score remains essentially unchanged at $Q_{max} = 0.658$, which is still within the high-risk band. This result highlights that the dominant risk shifts rather than disappears. Specifically, while the at-rest layer is significantly improved, the in-transit layer becomes the primary limiting factor for post-quantum readiness.

6.4. Comparative Analysis of Scenarios

The comparison between plaintext and protected scenarios reveals a structural redistribution of risk across layers. In the baseline configuration, the system is dominated by storage-related risk due to the absence of encryption and long confidentiality requirements. Once storage protection is introduced, this dominance is removed, and the risk profile becomes transport-driven. This shift is clearly illustrated in Figure 3, where the at-rest score decreases dramatically while the in-transit score remains unchanged.

The results confirm that improvements in one layer do not necessarily translate into overall risk elimination but instead reallocate the dominant source of exposure. Relative to the plaintext control, the protected condition reduces the weighted aggregate score Q_{agg} by 0.293 points, corresponding to a 41.5% reduction. At the layer level, the at-rest score Q_R falls by 0.734 points (80.4%), while Q_T and Q_U remain unchanged because the introduced control acts only on the storage path. The combination of aggregate and maximum-layer metrics is therefore necessary: the aggregate captures total urgency reduction, while Q_{max} shows that migration priority has shifted to–but has not been resolved at–the communication layer.

Within the scope of this case study, four observations follow from the layer-resolved scores: (i) the post-quantum exposure of the two reference configurations is markedly uneven across data states, and a single system-level metric does not surface this unevenness; (ii) evidence-based at-rest hardening produces a substantial aggregate risk reduction but does not eliminate system-level exposure; (iii) the transport layer remains the dominant source of post-quantum exposure in both scenarios because of its reliance on quantum-vulnerable public-key mechanisms; and (iv) layer-specific scoring is necessary to direct remediation effort to where it actually pays off. These observations are illustrative rather than statistically generalized; the latter would require a multi-host evaluation, as discussed in Section 7.6.

6.5. Sensitivity Analysis

This subsection reports a one-at-a-time sensitivity analysis of Q_{agg} and $Q_{\max}$ in the LUKS2-protected scenario. For each parameter family, the parameter is varied across a defensible range while all others are held at their nominal values from Section 3.9. The objective is twofold: to quantify how much the absolute score changes, and to test whether the qualitative ordering $\tilde{Q}_T > \tilde{Q}_U > \tilde{Q}_R$ in the protected scenario is preserved.

6.5.1. Sensitivity to the Quantum-Threat Horizon T_Q and Migration Buffer T_B

Because the temporal factors $r_{1,l}$ and $r_{2,l}$ scale inversely with T_Q and T_B, these two parameters have the largest mechanical effect on the score. Table 8 reports Q_{agg} and $Q_{\max}$ over $T_Q \in \{7, 10, 12, 15, 20\}$ years (representing aggressive-to-conservative quantum-threat scenarios consistent with the GRI 2024 timeline distribution [36]) and $T_B \in \{2, 3, 5\}$ years.

Table 8. Sensitivity of Q_{agg} (and $Q_{\max}$ in parentheses) to the quantum-threat horizon T_Q and migration buffer T_B in the LUKS2-protected scenario. Nominal point ($T_Q = 12$, $T_B = 3$) shown in bold. The transport dominance ordering $\tilde{Q}_T > \tilde{Q}_U > \tilde{Q}_R$ holds across the entire grid.

	$T_Q = 7$	$T_Q = 10$	$\boldsymbol{T_Q = 12}$	$T_Q = 15$	$T_Q = 20$
$T_B = 2$	0.477 (0.781)	0.450 (0.733)	0.439 (0.713)	0.427 (0.692)	0.416 (0.671)
$\boldsymbol{T_B = 3}$	0.453 (0.727)	0.426 (0.679)	**0.414 (0.658)**	0.402 (0.638)	0.391 (0.617)
$T_B = 5$	0.434 (0.687)	0.408 (0.640)	0.397 (0.620)	0.385 (0.600)	0.374 (0.580)

Three observations follow. First, the aggregate score Q_{agg} stays in the moderate band ($0.35 \leq Q < 0.65$) across the entire 5×3 grid. Second, $Q_{\max}$ stays in the high band ($Q \geq 0.65$) for short-horizon configurations ($T_Q \leq 12$) and crosses into the moderate band only for very conservative configurations ($T_Q \geq 15$ combined with $T_B \geq 3$). Third, the transport layer remains the dominant layer in every cell of the grid: the qualitative remediation conclusion is therefore robust to wide perturbations of the timeline assumptions.

6.5.2. Sensitivity to Attenuation Coefficients α and β

The Grover-bounded justification of $\alpha = 0.15$ and $\beta = 0.30$ in the subsection "Cryptographic Rationale for α and β" is qualitative rather than experimental. Table 9 reports Q_{agg} and $\tilde{Q}_R$ across a $\pm 50\%$ perturbation around each nominal coefficient.

Across the perturbation range, $\tilde{Q}_R$ varies between approximately 0.110 and 0.221, well within the low-risk band ($Q < 0.35$) for every tested combination. The transport layer ($\tilde{Q}_T = 0.658$) therefore continues to dominate $Q_{\max}$, and the aggregate Q_{agg} stays in the moderate band. The qualitative conclusion that LUKS2 protection drops the storage layer out of the dominant role is therefore not driven by the specific numerical choice of (α, β).

Table 9. Sensitivity of $\widetilde{Q}_R$ (LUKS2-protected) and Q_{agg} to the attenuation pair (α, β). Nominal point shown in bold.

(α, β)	$\widetilde{Q}_R$	Q_{agg}	Ordering Preserved?
$(0.075, 0.150)$	0.110	0.387	yes
$(0.150, 0.225)$	0.158	0.406	yes
$(\mathbf{0.150, 0.300})$	**0.179**	**0.414**	**yes**
$(0.150, 0.450)$	0.221	0.431	yes
$(0.225, 0.300)$	0.196	0.421	yes
$(0.300, 0.300)$	0.214	0.428	yes

6.5.3. Sensitivity to Aggregation Weights

As equal-weight aggregation may not reflect different business priorities (e.g., long-term storage versus temporary transit), Table 10 reports Q_{agg} under three alternative aggregation profiles. The transit-prioritized profile reflects organizations with high external attack-surface exposure; the storage-prioritized profile reflects archival-heavy environments such as legal record-keeping or genomic data repositories; and the equal-weight profile reflects organizations without a dominant priority.

Table 10. Sensitivity of Q_{agg} to alternative aggregation weights in the LUKS2-protected scenario.

Profile	(w_T, w_U, w_R)	Q_{agg}	Risk Band
Nominal (transit/at-rest balanced)	$(0.40, 0.20, 0.40)$	0.414	Moderate
Equal weights	$(1/3, 1/3, 1/3)$	0.411	Moderate
Transit-prioritized	$(0.60, 0.20, 0.20)$	0.510	Moderate
Storage-prioritized (long retention)	$(0.20, 0.20, 0.60)$	0.318	Low
In-use heavy (runtime-critical)	$(0.30, 0.40, 0.30)$	0.409	Moderate

The aggregate score varies between 0.318 (storage-prioritized) and 0.510 (transit-prioritized), reflecting the genuinely different risk pictures faced by organizations with different strategic priorities. Importantly, the transport-prioritized profile leaves the score firmly in the moderate band, while the storage-prioritized profile pushes the score into the low band only because the storage layer has been hardened. In every profile, the transport layer remains the dominant single layer ($Q_{\max} = 0.658$), and so the migration prioritization conclusion does not depend on the choice of aggregation weights.

6.5.4. Sensitivity to Risk Thresholds

Moving the moderate/high threshold from 0.65 to 0.60 would re-classify two of the four sensitivity cells in Table 8 as high-risk; moving it to 0.70 would re-classify the nominal $Q_{\max}$ cell as moderate. The bands are therefore advisory rather than absolute, and operators with stronger or weaker risk appetites should adjust accordingly. The 0.35/0.65 partition adopted in this paper is consistent with the three-band partitioning recommended by CARAF [25].

6.6. Positioning of QARS Relative to Existing Frameworks

The purpose of this subsection is to position QARS relative to existing post-quantum migration and crypto-agility frameworks, not to benchmark its numerical performance against them. Of the frameworks compared below, only CARAF [25] produces a structured risk output, and that output is ordinal and system-level rather than continuous and layer-resolved; the remaining frameworks deliver migration playbooks, discovery artifacts, asset inventory schemas, or wire-format specifications rather than risk scores. A direct numerical

comparison on the same host evidence is therefore not meaningful, and Table 11 accordingly characterizes each framework along the dimensions that define its scope output.

The positioning clarifies the complementary niche of QARS rather than asserting superiority: QARS provides a numerical, layer-resolved score that can serve as input to higher-level frameworks rather than as a replacement for them. CARAF's roadmap stage, ETSI's stage-2 migration plan, and NIST's interoperability testing all benefit from a quantitative layer-resolved input that QARS produces. Conversely, QARS does not replace cryptographic discovery (NIST SP 1800-38B), wire-format specification (IETF), or asset inventory schema (CBOM); it consumes those artifacts.

Table 11. Comparison of QARS with leading post-quantum migration and risk assessment frameworks. "Yes", "Partial", and "No" indicate whether the framework supports the listed feature.

Framework	Input	Output	Layer-Resolved?	Symmetric vs. Shor Distinction?	Operational Output
Mosca timeline argument [28]	Time horizons X, Y, Z	Boolean ($X + Y > Z$)	No	Implicit (asymmetric only)	Strategic alarm
CARAF [25]	Threat vector + asset inventory + asset value (qualitative)	5-step roadmap; ordinal categories	No (asset-by-asset)	No	Mitigation roadmap
ETSI TR 103 619/ / TR 104 016 [8,9]	Cryptographic asset inventory	Stage-gated migration plan	No (system-wide)	Partial (separates QSC vs. non-QSC)	Repeatable migration framework
NIST SP 1800-38 [7]	Cryptographic discovery scans	Discovery and interoperability artifacts	No	No	Discovery and testing playbook
CycloneDX CBOM [27]	Crypto asset enumeration (algorithm, key, certificate)	Structured BOM (JSON/XML)	No	Encoded as asset metadata	Cryptographic inventory
IETF hybrid TLS [47,48]	Protocol parameters	Hybrid key-exchange profile	No (transport only)	Yes (within transport)	Wire-format specification
QARS (this work)	Host-level evidence (scans, configs, FDE state)	Layer-resolved $\widetilde{Q}_l$ + dual-metric (Q_{agg}, $Q_{\max}$)	**Yes (transit/ in-use/at-rest)**	**Yes (Grover-bounded α, β)**	Numerical layer ranking + remediation ordering

Bold entries indicate the framework proposed in this work (QARS).

7. Discussion

The results demonstrate that post-quantum risk in enterprise infrastructure is inherently uneven across data layers and cannot be accurately represented by a uniform system-level metric. Instead, the proposed layer-specific approach reveals how different protection mechanisms influence both the magnitude and distribution of risk.

7.1. Interpretation of Layer-Specific Risk

The experimental findings confirm that the dominant source of post-quantum risk depends on the protection state of the infrastructure. In the plaintext control scenario, the at-rest layer exhibits the highest score due to the combination of long confidentiality requirements and the absence of cryptographic protection. This aligns with theoretical expectations, as unprotected long-term storage is highly vulnerable to delayed-decryption attacks in a post-quantum context.

However, once storage protection is introduced, the risk profile changes significantly. The at-rest layer no longer dominates due to the application of symmetric encryption and the associated attenuation of quantum vulnerability. Instead, the transport layer becomes the primary source of risk. This shift highlights a key property of post-quantum migration: risk is not eliminated by improving a single layer, but is rather redistributed across the system.

The divergence between the aggregate score and the maximum-layer score provides important insight into infrastructure security assessment. The reduction in the weighted aggregate score from high to moderate indicates that storage hardening has a substantial

impact on overall system posture. At the same time, the persistence of a high maximum-layer score reveals that critical vulnerabilities remain unresolved.

This distinction directly addresses RQ1, demonstrating that layer-specific scoring provides a more informative and operationally relevant representation of post-quantum exposure than uniform system-level assessment. The results further address RQ2 by showing that strengthening the at-rest layer significantly reduces aggregate risk, while RQ3 is answered by confirming that transport exposure remains the dominant migration priority even after storage protection is applied.

From an operational perspective, this implies that organizations should not interpret improvements in aggregate metrics as sufficient evidence of readiness. Instead, remediation efforts must prioritize the most exposed layer, which in this case remains the communication infrastructure.

7.2. Implications for Post-Quantum Migration Strategy

The findings suggest that post-quantum migration should follow a layer-aware prioritization strategy. While storage encryption is effective in mitigating long-term confidentiality risks, it does not address vulnerabilities associated with public-key-based communication protocols. Consequently, transport-layer modernization, including the adoption of quantum-resistant key exchange and authentication mechanisms, should be treated as the primary migration objective.

At the same time, the moderate score of the in-use layer indicates that runtime controls can effectively reduce exposure without requiring immediate cryptographic transformation. This supports a phased migration approach in which transport security is addressed first, followed by gradual improvements in runtime and storage environments.

7.3. Operational Transport-Layer Modernization Plan

The persistently high Q_T in both scenarios is a direct consequence of three observable evidence types in `Hosts_info.xlsx`: the presence of a TLS 1.0 listener, the negotiation of CBC-mode cipher suites, and the use of an RSA-only certificate without hybrid key-exchange support. Each of these is independently remediable, and the QARS framework provides explicit hooks for re-evaluating the score after remediation. The recommended four-stage remediation sequence is therefore as follows.

Stage 1–Removal of obsolete TLS versions. The TLS 1.0 endpoint is disabled at the load-balancer and on each backend service; the supported version floor is raised to TLS 1.2. This change does not, on its own, alter the layer's quantum exposure α, but it removes a non-quantum cryptographic weakness that would otherwise dominate the residual risk and obscure the quantum signal during follow-up audits. The corresponding cell in `Hosts_info.xlsx` (sheet "In-Transit", field "TLS versions") is updated, and the transit-layer score is recomputed.

Stage 2–Retirement of weak-cipher suites. CBC-mode and SHA-1-based suites are removed, and the AEAD-only profile (AES-GCM, ChaCha20-Poly1305) is enforced. This is again a non-quantum hardening step, but it is required before the hybrid key-exchange experiments in Stage 3 can produce meaningful results because hybrid suites are only standardized on top of AEAD profiles [47,48].

Stage 3–Baselining TLS 1.3. TLS 1.3 is made the default and the minimum negotiable version. At this point, the transport layer carries no exploitable confidentiality weakness against a classical adversary, but it remains fully quantum-vulnerable: the X25519 and ECDHE key exchanges are recoverable by Shor's algorithm [1,10]. The QARS quantum-vulnerability flag for the transit layer therefore stays at $\alpha = 1$, and Q_T does not yet decrease.

Stage 4–Hybrid post-quantum key exchange. The deployment is upgraded to a hybrid key-exchange group such as `X25519MLKEM768`, in line with the IETF hybrid TLS design [47,48] and the NIST migration testbeds [7]. Once the hybrid group is negotiated for at least 95% of inbound traffic and signing certificates remain classical, the transit-layer attenuation can be reduced from $\alpha = 1$ to the at-rest setting $\alpha = 0.15$, with the precise value determined by the operator's risk tolerance and the proportion of pre-quantum-only fallbacks observed in monitoring. The recomputed Q_T is then re-aggregated with the existing Q_U and Q_R to produce the post-remediation Q_{agg} and $Q_{\max}$.

This staged plan illustrates a key methodological property of QARS: the score is not a one-shot diagnosis but an instrument that can be recomputed after each remediation step, enabling operators to verify that hardening is actually moving the metric in the expected direction. The same recomputation hook applies to in-use and at-rest changes; what is specific to transport is the asymmetry between non-quantum hardening (Stages 1–2) and quantum hardening (Stages 3–4), and the fact that only the latter modifies α.

7.4. Saturation Behavior of the Bounded Operators

The model contains two bounded operators: the $\min(\cdot, 2)$ cap on temporal factors r_1 and r_2 in Equation (2), and the $\text{clip}_{[0,1]}$ operator on the layer score in Equation (1). Neither operator activates in the case-study data reported in Section 6. The temporal factors range from 0.167 to 1.250 (Table 7), well below the cap of 2; the layer scores range from 0.179 to 0.913, well below the clip threshold of 1. Under the nominal calibration ($T_Q = 12$, $T_B = 3$), the temporal cap would activate only for $T^{conf} > 24$ years or $T^{mig} > 6$ years, and the layer-score clip would activate only when the weighted sum of the factors exceeds unity. These conditions describe archival storage with very long retention horizons or estates with stalled migration programs, neither of which is present in the evaluated configurations.

The ceiling-effect compression discussed in earlier drafts of this paper–in which two infrastructures with materially different underlying exposures receive the same saturated score–therefore describes a known limitation of the bounded-aggregation design that is not exercised by the present case study. It is recorded as a limitation in Section 7.6.

A symmetric floor concern operates at the lower end. After at-rest hardening, the at-rest layer score drops to 0.179, well above zero but below the moderate threshold of 0.35. This score does not reflect residual operational risk from key-management practices, vault-unsealing procedures, or backup encryption coverage; these are tracked separately through in-use evidence. Operators interpreting low at-rest scores should treat them as conditional on sound key management rather than as absolute statements of safety.

7.5. Cross-Layer Dependencies and Their Effect on Remediation Order

The cross-layer dependencies introduced in Section 3.11 have direct implications for how remediation should be sequenced in practice. Three sequencing rules follow from these cross-layer mappings. First, transport hardening must precede any reliance on encrypted application-layer payloads as the primary in-use control: as long as the TLS endpoint terminates plaintext traffic on a host with high in-use exposure (active SUID binaries, shell access, weak process isolation), the assumption that the application sees only ciphertext is unsafe, and double-counting of `stunnel`-style overlays should be avoided per subsection "Classification of Redis-Related Evidence". Second, in-use hardening must precede any claim that at-rest encryption protects against runtime key extraction: if the in-use layer contains conditions that allow a privileged process to read the LUKS2 master key from the kernel memory, the at-rest score is operationally meaningless. The model captures this by treating in-use evidence as a precondition in the provenance table, but the score

itself does not enforce ordering. Third, key-management workflows for at-rest encryption (vault unsealing, key rotation, escrow) sit at the boundary between in-use and at-rest, and improvements in either layer can shift the effective α of the other: a hardware-isolated key store reduces the residual exposure of LUKS2 even when the in-use layer is otherwise unchanged. These observations motivate the layered remediation order "transport $\rightarrow$ in-use $\rightarrow$ at-rest key management" and explain why the maximum-layer score $Q_{\max}$ remains a more reliable migration priority indicator than Q_{agg} alone.

7.6. Limitations

Several limitations of the present study should be acknowledged.

The evaluation reported in this paper is an illustrative case study based on two hosts in a within-case comparative design. This supports the demonstration of the evidence-to-score pipeline and a controlled before/after comparison of storage hardening, but it does not constitute a sample from which statistical inference can be drawn: this case study does not establish a distribution of layer scores across independently configured hosts, does not quantify the variance of the observed layer ordering, and does not validate fleet-level aggregation rules. A multi-host extension across independently configured hosts, along with reported per-layer score ranges, is identified as primary direction for future work. The claims of this paper should be read as qualitative and configuration-specific rather than as statistical properties of QARS on a production fleet.

The cross-layer interdependencies discussed in Section 7.5 are reasoned about analytically and via the provenance taxonomy, but they are not exercised under load. A multi-host deployment with realistic east–west traffic, shared key-management infrastructure, and observability pipelines would allow these dependencies to be measured directly rather than asserted. This is a known limitation that constrains the strength of conclusions about emergent system behavior.

The mapping from observed evidence to scoring parameters involves heuristic transformations that, although transparent and auditable, may not fully capture all aspects of real-world risk. For example, sensitivity and exposure are derived from selected indicators and may vary depending on domain-specific factors not included in the model.

The quantum threat horizon T_Q and migration time T_B are calibrated against published estimates, but as the sensitivity analysis in Section 6.5 shows, varying them within the credible range moves the protected Q_{agg} by approximately ± 0.05. Operators with strong priors on either parameter should re-run the recomputation notebook with their own values rather than treat the reported numbers as definitive.

The attenuation model uses two discrete settings for α (0.15 for symmetric protection, and 1.0 for unprotected or public-key-only paths) and a single setting for β. While this is appropriate for symmetric encryption versus public-key cryptography, hybrid and emerging cryptographic schemes (lattice-based KEMs combined with classical ECDHE, hash-based signatures alongside RSA) occupy intermediate positions on the quantum-vulnerability spectrum and would benefit from a continuous α parameterized by residual classical strength and Grover-bounded symmetric strength. This is a planned direction for follow-up work.

The bounded operators discussed in Section 7.4 would reduce discriminative power in the environments reaching their saturation regions; these regions are not reached in the present case study but constrain use of the raw score in archival or migration-stalled estates.

8. Conclusions

This paper presented an evidence-based Quantum-Adjusted Risk Scoring (QARS) model for evaluating post-quantum exposure across enterprise infrastructure. The proposed approach addresses a key limitation of existing methods by distinguishing between data in transit, in use, and at rest, as well as by enabling direct comparison of risk across these layers within a unified framework and by deriving every input from observable host evidence rather than from analyst opinion.

The contribution sits in a specific niche relative to prior work: the model is layer-resolved (where Mosca's inequality is system-level), it produces a continuous numerical score (where CARAF and ETSI TR 104 016 produce roadmap stages), and it consumes the cryptographic discovery artifacts standardized by NIST SP 1800-38B and CycloneDX CBOM rather than re-implementing them. Three components are reused from prior work–the time-budget intuition of Mosca, the qualitative ordering Shor $\gg$ Grover used in attenuation, and the layer-resolved framing inherited from the data-state taxonomy–and three are introduced here: a Grover-bounded attenuation factor with explicit cryptographic justification, an evidence-to-feature mapping that makes every input auditable, and a dual-metric (Q_{agg}, $Q_{\max}$) reporting scheme that separates aggregate progress from migration priority.

The case study results illustrate that post-quantum risk can be markedly uneven across data states in realistic configurations and that this unevenness is not visible under uniform system-level assessment. Whether this pattern holds across a representative production estate is an open question that requires a multi-host evaluation; the present work is restricted to a two-host illustrative scope and reports parameter sensitivity rather than fleet-level statistics. At the design level, the introduction of quantum-vulnerability attenuation for symmetric protection mechanisms significantly alters the interpretation of storage-related risk, preventing overestimation of urgency in protected at-rest environments. At the same time, the findings show that improvements in one layer do not eliminate overall risk but instead shift the dominant source of exposure: in the evaluated scenarios, storage hardening reduces the aggregate score by 41.5% (from 0.707 to 0.414), yet the transport layer remains the primary bottleneck at $Q_T = 0.658$ due to its dependence on quantum-vulnerable public-key mechanisms.

The sensitivity analysis confirms that this qualitative ordering is robust to plausible perturbations of the calibration parameters. Across a 5×3 grid of T_Q and T_B values, Q_{agg} in the protected scenario varies between 0.374 and 0.477, and $Q_{\max}$ remains transport-dominated for all settings with $T_Q \leq 12$ years. Perturbing α and β by $\pm 50\%$ shifts the protected Q_{agg} by less than 0.05 points. Alternative aggregation profiles that prioritize transit, storage, or in-use exposure produce Q_{agg} values in the range $[0.318, 0.510]$ but never alter the layer ordering. The qualitative conclusion that transport modernization is the priority therefore does not depend on the specific defaults chosen.

The four-stage operational transport-layer modernization plan–TLS 1.0 retirement, weak-cipher removal, TLS 1.3 baselining, and hybrid post-quantum key exchange–shows how the QARS score can be re-evaluated after each remediation step using the same observation pipeline, providing operators with a continuous progress indicator rather than a one-shot diagnosis.

The dual-metric evaluation, combining weighted aggregate and maximum-layer scores, provides complementary perspectives for decision-making. While the aggregate score reflects overall system posture, the maximum-layer score identifies the most critical component requiring immediate remediation. Cross-layer interdependencies–in-use compromise leaking at-rest keys, transit compromise propagating to in-use, and at-rest compromise unsealing transport credentials–imply that the dominant-layer priority should always be implemented together with hardening of the layers that gate access to its key material.

https://doi.org/10.3390/electronics15122546

The proposed model is based on observable host evidence, ensuring that risk assessments remain transparent, reproducible, and aligned with real-world infrastructure conditions; the anonymized dataset, recomputation notebook, and evidence-collection commands are scheduled for public release alongside this paper. Limitations include the two-host validation scope, the binary attenuation parameterization, and the discriminative compression that the bounded operators would introduce in environments, reaching the saturation regions described in Section 7.4. Future work should extend the model to larger and more heterogeneous environments with measured cross-layer traffic, replace the binary α with a continuous Grover-bounded function, refine parameter estimation using empirical datasets, and incorporate additional dimensions such as integrity and availability. Integration of hybrid and evolving post-quantum cryptographic mechanisms–particularly emerging hash-based and code-based signatures–into the scoring framework represents an important direction for further research.

The QARS model provides a practical and extensible approach for assessing post-quantum readiness, supporting informed decision-making and enabling structured prioritization of migration efforts across enterprise infrastructure.

Author Contributions: Conceptualization, Š.G., S.K. and R.B.; methodology, Š.G. and R.B.; software, S.K. and M.C.; validation, Š.G., S.K. and R.B.; formal analysis, Š.G. and R.B.; investigation, Š.G., S.K. and R.B.; resources, S.K. and M.C.; data curation, S.K. and R.B.; writing–original draft preparation, Š.G., S.K. and R.B.; writing–review and editing, R.B. and S.K.; visualization, S.K. and R.B.; supervision, Š.G.; project administration, Š.G.; funding acquisition, Š.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research work was conducted as part of the execution of Project "Pionier-Q interconnection with the EuroQCI's space segment (24-EU-DIG-PIONIER-Q-SAT)", Project no.: 101249721, funded by European Union under The Connecting Europe Facility (CEF) in Digital program.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: An anonymized version of the four-sheet workbook used in this study (`Hosts_info_anon.xlsx`), the Python 3.11 recomputation notebook `qars_recompute.ipynb`), and the evidence-collection commands listed in Section 5.4 are scheduled for release in a public Git repository at the time of publication, in compliance with the host institution's data sovereignty and cybersecurity policies. Until the repository is published, these artifacts are available from the corresponding author upon reasonable request. Identifying information (host names, IP addresses, certificate fingerprints) has been redacted from the released material; this redaction does not affect the reproducibility of the reported scores.

Acknowledgments: During the preparation of this manuscript, the authors used a generative AI tool (ChatGPT 5.4) for the revision of the English language. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Shor, P.W. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *Siam J. Comput.* **1997**, *26*, 1484–1509. [CrossRef]
2. Grover, L.K. A Fast Quantum Mechanical Algorithm for Database Search. In *Proceedings of 28th Annual ACM Symposium on Theory of Computing (STOC '96)*; ACM: New York, NY, USA, 1996; pp. 212–219. [CrossRef]
3. *FIPS 203*; Module-Lattice-Based Key-Encapsulation Mechanism Standard. National Institute of Standards and Technology: Gaithersburg, MD, USA, 2024. Available online: https://csrc.nist.gov/pubs/fips/203/final (accessed on 25 March 2026).
4. *FIPS 205*; Stateless Hash-Based Digital Signature Standard. National Institute of Standards and Technology: Gaithersburg, MD, USA, 2024. Available online: https://csrc.nist.gov/pubs/fips/205/final (accessed on 25 March 2026).

5. *FIPS 204*; Module-Lattice-Based Digital Signature Standard. National Institute of Standards and Technology: Gaithersburg, MD, USA, 2024. Available online: https://csrc.nist.gov/pubs/fips/204/final (accessed on 25 March 2026).
6. National Institute of Standards and Technology. *Post-Quantum Cryptography*; NIST Computer Security Resource Center Project Page; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2025. Available online: https://csrc.nist.gov/projects/post-quantum-cryptography (accessed on 25 March 2026).
7. National Institute of Standards and Technology. *Migration to Post-Quantum Cryptography*; Technical Report NIST SP 1800-38 (Volumes A, B, C); Preliminary Draft Practice Guide; NIST National Cybersecurity Center of Excellence: Rockville, MD, USA, 2023.
8. European Telecommunications Standards Institute. *CYBER; Quantum-Safe Cryptography (QSC); A Repeatable Framework for Quantum-Safe Migrations*; Technical Report ETSI TR 104 016 V1.1.1; ETSI: Valbonne, France, 2024.
9. European Telecommunications Standards Institute. *CYBER; Migration Strategies and Recommendations to Quantum Safe Schemes*; Technical Report ETSI TR 103 619 V1.1.1; ETSI: Valbonne, France, 2020.
10. National Security Agency. *Commercial National Security Algorithm Suite 2.0 (CNSA 2.0)*; Technical Report, NSA Cybersecurity Advisory; NSA: Fort George G. Meade, MD, USA, 2022.
11. Barker, E.; Roginsky, A. *Transitioning the Use of Cryptographic Algorithms and Key Lengths*; NIST Special Publication 800-131A Rev. 2; NIST: Gaithersburg, MD, USA, 2019. [CrossRef]
12. UK National Cyber Security Centre. *Timelines for Migration to Post-Quantum Cryptography*; NCSC Guidance; UK National Cyber Security Centre: London, UK, 2025. Available online: https://www.ncsc.gov.uk/guidance/pqc-migration-timelines (accessed on 30 April 2026).
13. Joseph, D.; Misoczki, R.; Manzano, M.; Tricot, J.; Pinuaga, F.D.; Lacombe, O.; Leichenauer, S.; Hidary, J.; Venables, P.; Hansen, R. Transitioning organizations to post-quantum cryptography. *Nature* **2022**, *605*, 237–243. [CrossRef] [PubMed]
14. Von Nethen, N.; Wiesmaier, A.; Alnahawi, N.; Henrich, J. PMMP-PQC Migration Management Process. In Proceedings of the 2024 European Interdisciplinary Cybersecurity Conference, Xanthi, Greece, 5–6 June 2024; pp. 144–154.
15. Malina, L.; Dobias, P.; Hajny, J.; Choo, K.K.R. On deploying quantum-resistant cybersecurity in intelligent infrastructures. In Proceedings of the 18th International Conference on Availability, Reliability and Security, Benevento, Italy, 29 August–1 September 2023; pp. 1–10.
16. Ylonen, T.; Lonvick, C. (Eds.) *RFC 4253: The Secure Shell (SSH) Transport Layer Protocol*; RFC: Wilmington, DE, USA, 2006. [CrossRef]
17. Rescorla, E. *RFC 8446: The Transport Layer Security (TLS) Protocol Version 1.3*; RFC: Wilmington, DE, USA, 2018. [CrossRef]
18. Sikeridis, D.; Kampanakis, P.; Devetsikiotis, M. Assessing the overhead of post-quantum cryptography in TLS 1.3 and SSH. In Proceedings of the 16th International Conference on Emerging Networking EXperiments and Technologies, Barcelona, Spain, 1–4 December 2020; pp. 149–156.
19. Sosnowski, M.; Wiedner, F.; Hauser, E.; Steger, L.; Schoinianakis, D.; Gallenmüller, S.; Carle, G. The performance of post-quantum tls 1.3. In Proceedings of the Companion of the 19th International Conference on Emerging Networking Experiments and Technologies, Paris, France, 5–8 December 2023; pp. 19–27.
20. Sikeridis, D.; Kampanakis, P.; Devetsikiotis, M. Post-Quantum Authentication in TLS 1.3: A Performance Study. In Proceedings of the 27th Annual Network and Distributed System Security Symposium (NDSS 2020), San Diego, CA, USA, 23–26 February 2020.
21. Redis Ltd. TLS. Official Redis Documentation. 2026. Available online: https://redis.io/docs/latest/operate/oss_and_stack/management/security/encryption/ (accessed on 25 March 2026).
22. Redis Ltd. Redis Persistence. Official Documentation. 2026. Available online: https://redis.io/docs/latest/operate/oss_and_stack/management/persistence/ (accessed on 25 March 2026).
23. Popa, R.A. Confidential computing or cryptographic computing? Tradeoffs between cryptography and hardware enclaves. *Queue* **2024**, *22*, 108–132. [CrossRef]
24. Feng, D.; Qin, Y.; Feng, W.; Li, W.; Shang, K.; Ma, H. Survey of research on confidential computing. *IET Commun.* **2024**, *18*, 535–556. [CrossRef]
25. Ma, C.; Colon, L.; Dera, J.; Rashidi, B.; Garg, V. CARAF: Crypto Agility Risk Assessment Framework. *J. Cybersecur.* **2021**, *7*, tyab013. [CrossRef]
26. Hohm, J.; Heinemann, A.; Wiesmaier, A. Towards a maturity model for crypto-agility assessment. In *International Symposium on Foundations and Practice of Security*; Springer: Berlin/Heidelberg, Germany, 2022; pp. 104–119.
27. OWASP CycloneDX. *Cryptography Bill of Materials (CBOM): Authoritative Guide*; CycloneDX Specification v1.6, ECMA-424; OWASP Foundation: Wilmington, DE, USA, 2024. Available online: https://cyclonedx.org/guides/OWASP_CycloneDX-Authoritative-Guide-to-CBOM-en.pdf (accessed on 30 April 2026).
28. Mosca, M. Cybersecurity in an era with quantum computers: Will we be ready? *IEEE Secur. Priv.* **2018**, *16*, 38–41. [CrossRef]
29. Howison, M.; Angell, M.; Hastings, J.S. Protecting sensitive data with secure data enclaves. *Digit. Gov. Res. Pract.* **2024**, *5*, 1–11. [CrossRef]

30. European Parliament and Council. Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2 Directive). *Off. J. Eur. Union* **2022**, 80–152. Available online: https://eur-lex.europa.eu/eli/dir/2022/2555/oj (accessed on 12 March 2026).
31. Carvalho, O.; Apolinário, F.; Escravana, N.; Ribeiro, C. CIIA: Critical infrastructure impact assessment. In Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing, Virtual, 25–29 April 2022; pp. 124–132.
32. Segovia-Ferreira, M.; Rubio-Hernan, J.; Cavalli, A.; Garcia-Alfaro, J. A survey on cyber-resilience approaches for cyber-physical systems. *ACM Comput. Surv.* **2024**, *56*, 1–37. [CrossRef]
33. Ding, J.; Atif, Y.; Andler, S.F.; Lindström, B.; Jeusfeld, M. CPS-based threat modeling for critical infrastructure protection. In *ACM SIGMETRICS Performance Evaluation Review*; ACM: New York, NY, USA, 2017; Volume 45, pp. 129–132.
34. Gallardo, A.; Erbes, R.; Le Blanc, K.; Bauer, L.; Cranor, L.F. Interdisciplinary approaches to cybervulnerability impact assessment for energy critical infrastructure. In Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, 11–16 May 2024; pp. 1–24.
35. Grigaliūnas, Š.; Brūzgienė, R. Towards a Unified Quantum Risk Assessment. *Electronics* **2025**, *14*, 3338. [CrossRef]
36. Mosca, M.; Piani, M. Quantum Threat Timeline Report 2024. In *Global Risk Institute Report*; Global Risk Institute: Toronto, ON, Canada, 2024.
37. Alnahawi, N.; Schmitt, N.; Wiesmaier, A.; Heinemann, A.; Graßmeyer, T. *On the State of Crypto Agility*; BSI: Bonn, Germany, 2022.
38. Bindel, N.; Herath, U.; McKague, M.; Stebila, D. Transitioning to a quantum-resistant public key infrastructure. In *Proceedings of the International Workshop on Post-Quantum Cryptography*; Springer: Berlin/Heidelberg, Germany, 2017; pp. 384–405.
39. Bürstinghaus-Steinbach, K.; Krauß, C.; Niederhagen, R.; Schneider, M. Post-quantum TLS on embedded systems: Integrating and evaluating Kyber and SPHINCS+ with mbed TLS. In Proceedings of the 15th ACM Asia Conference on Computer and Communications Security, Taipei, Taiwan, 5–9 October 2020; pp. 841–852. [CrossRef]
40. Tasopoulos, G.; Dimopoulos, C.; Fournaris, A.P.; Zhao, R.K.; Sakzad, A.; Steinfeld, R. Energy consumption evaluation of post-quantum TLS 1.3 for resource-constrained embedded devices. In Proceedings of the 20th ACM International Conference on Computing Frontiers, Bologna, Italy, 9–11 May 2023; pp. 366–374.
41. Paul, S.; Kuzovkova, Y.; Lahr, N.; Niederhagen, R. Mixed certificate chains for the transition to post-quantum authentication in TLS 1.3. In Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security, Nagasaki, Japan, 30 May–2 June 2022; pp. 727–740.
42. Sikeridis, D.; Huntley, S.; Ott, D.; Devetsikiotis, M. Intermediate certificate suppression in post-quantum TLS: An approximate membership querying approach. In Proceedings of the 18th International Conference on Emerging Networking EXperiments and Technologies, Rome, Italy, 6–9 December 2022; pp. 35–42.
43. Paul, S.; Schick, F.; Seedorf, J. TPM-Based Post-Quantum Cryptography: A Case Study on Quantum-Resistant and Mutually Authenticated TLS for IoT Environments. In *Proceedings of the 16th International Conference on Availability, Reliability and Security (ARES 2021)*; ACM: New York, NY, USA, 2021. [CrossRef]
44. Kwiatkowski, K.; Sullivan, N.; Langley, A.; Levin, D.; Mislove, A. Measuring TLS Key Exchange with Post-Quantum KEM. In Proceedings of the Second PQC Standardization Conference, Santa Barbara, CA, USA, 22–24 August 2019.
45. Mo, F.; Tarkhani, Z.; Haddadi, H. Machine learning with confidential computing: A systematization of knowledge. *ACM Comput. Surv.* **2024**, *56*, 1–40. [CrossRef] [PubMed]
46. Diesburg, S.M.; Wang, A.I.A. A survey of confidential data storage and deletion methods. *ACM Comput. Surv.* **2010**, *43*, 1–37. [CrossRef]
47. Stebila, D.; Fluhrer, S.; Gueron, S. Hybrid Key Exchange in TLS 1.3. Internet-Draft draft-ietf-tls-hybrid-design, Internet Engineering Task Force (IETF), 2024. Available online: https://datatracker.ietf.org/doc/draft-ietf-tls-hybrid-design/ (accessed on 15 March 2026).
48. Hale, B.; Reddy, T.; Driscoll, F.; Banerjee, N.; Kampanakis, P. *RFC 9794: Terminology for Post-Quantum Traditional Hybrid Schemes*; RFC: Wilmington, DE, USA, 2025. [CrossRef]