SETSCI Conference Proceedings
25, 42–47, 2026

International Conference on Open Finance
September 24, 2026, Boston, MA, USA

https://doi.org/10.36287/setsci.25.8.001
2687-5527/ © 2026 The Authors. Published by SETSCI

Quantum-Ready Intelligent Digital Systems for Open Finance: Integrating Post-Quantum Security and Evidence-Gated Quantum Computing

Ashwini Kumar Rath
Batoi Research, Bhubaneswar, India (ashwini.rath@batoi.com)

Abstract – Open-finance services must prepare for cryptographic transitions while evaluating quantum-computing opportunities against classical alternatives. These decisions require different evidence but concern shared data, software, suppliers, and owners. This paper develops a design-science framework through a multivocal review of 60 sources. It links a system evidence graph, trust-transition register, quantum-experiment register, and capability-and-decision register across seven technical and organizational layers. The framework separates preparation priority from deployment feasibility, constrains maturity judgments by evidence, and gates experiments through authority, baseline quality, full-workflow accounting, and stopping rules. A hypothetical bank-account aggregation service illustrates discovery, a maturity cap, and an experiment stop decision. A change-event trace identifies the proposed integration beyond established migration records without claiming measured superiority. Two weight scenarios reverse the priority order of two work packages; criterion-removal and input sensitivity checks expose the dependence on exposure assumptions. The contribution is a specified linkage between migration, experimentation, and accountable review. No participants or operational deployment were assessed. The framework remains unvalidated; the example demonstrates inspectable decision mechanics, not security assurance or decision utility.

Keywords – quantum readiness; open finance; post-quantum cryptography; crypto-agility; intelligent digital systems; quantum security; hybrid quantum computing; digital trust

I. Introduction

Open finance extends customer-permissioned data access from banking into savings, credit, insurance, pensions, and investments [1]. The Basel Committee identifies cybersecurity and third-party challenges in data sharing [2]; BIS work emphasizes secure identification, authentication, and API interoperability [3]. These arrangements distribute ownership of consent, identities, suppliers, and audit evidence across institutions. A service owner may depend on another institution's migration or verification profile without controlling it. Quantum technology affects this environment through two pathways.

The first pathway is computational opportunity. Shor established polynomial-time quantum algorithms for factoring and discrete logarithms [4]; Grover established a quadratic oracle-query improvement for unstructured search [5]. Neither predicts commercial usefulness or machine availability. Experimental claims also depend on their comparator: Google's Sycamore sampling experiment [6] was followed by a substantially faster classical approach to that sampling family [7]. Noisy intermediate-scale algorithms span heterogeneous tasks, hardware, and error conditions [8]. Application proposals therefore require current, problem-specific classical comparisons.

The second pathway is digital trust. A sufficiently capable fault-tolerant quantum computer would threaten RSA, finite-field Diffie-Hellman, and elliptic-curve systems [4]. Harvest-now-decrypt-later exposure concerns captured ciphertext whose confidentiality must endure; signatures require separate authenticity and verification-lifetime analysis. The Mosca timing model relates protection lifetime and full migration lead time to the uncertain arrival of a cryptographically relevant quantum computer [9]. NIST finalized ML-KEM, ML-DSA, and SLH-DSA in 2024 [10]–[12] and subsequently issued KEM guidance [13]. Standards do not replace dependency discovery, protocol redesign, supplier coordination, rollback testing, or continuing errata checks.

Security preparation should be anticipatory; computational investment should be problem-led and gated by current classical baselines. A single adoption score obscures their different evidence requirements.

This paper asks how an open-finance service can govern both pathways without conflating theory, experiments, standards, and aspirations. It proposes four linked operational records for one intelligent digital system with a defined scope. The integration proposition is that shared evidence references and explicit review triggers can make the consequences of one change inspectable across migration, experiment, and capability decisions. The paper specifies this mechanism and illustrates it; improved decision utility remains a hypothesis.

II. Research Approach

The multivocal review used Google Scholar and web discovery of publisher records, plus direct official standards and government searches, on August 1–2, 2026; open-finance sources were added on August 15–16. English-language material from 2018–2026 was prioritized, with seminal exceptions. Queries covered readiness, PQC migration, crypto-agility, advantage and classical responses, hybrid workflows, governance, QKD, and Indian policy. Inclusion required relevance to a research question or design requirement; marketing was not scientific evidence. QKD was searched for context but its deployment is outside this artifact's scope: PQC migration and quantum-computing experiments are distinct from quantum key distribution [9].

The ledger contains 60 sources; the 29 references are a citation subset, not another screening stage. Records retain provenance, claims, limitations, access status, and verification dates. Inspection ranged from full text to abstracts or metadata, departing from the initial full-text eligibility rule. Claims exclude uninspected numerical and implementation findings. Authenticated database coverage, complete candidate-level exclusion counts, and independent screening are undocumented; no systematic-review claim is made. A September 5 currency check updated existing standards/protocol records without expanding the corpus.

Synthesis translated inspected claims into requirements and then records, caps, and gates. NIST supplies cross-layer crypto-agility guidance [14]. ETSI already includes identifiers, dependency digraphs, external dependencies, and priority analysis [15], with separate hybrid-deployment guidance [16]. Readiness studies address capability and staged experimentation [17], [18]. Table 2 traces these inherited requirements to the four proposed records; the linking rules and numerical conventions are author-designed, not validated source findings. Supporting records use separate evidence-matrix and claim-register namespaces with a semantic crosswalk; source access corrections remain dated.

Design science distinguishes artifact construction from evaluation [19], [20]. This paper reports construction and a synthetic walkthrough with arithmetic stress tests, not the planned empirical evaluation. No participants or confidential organizational data were involved.

III. Synthesis Results and Design Requirements

A. Threat and opportunity must remain distinct

Grover's query-complexity improvement depends in practice on oracle construction, depth, parallelization, error correction, and hardware [5]. Unlike Shor-vulnerable public-key systems, symmetric ciphers face a quadratic generic key-search speedup, not the same structural break [4], [5]. Key size, implementation, and protection lifetime still require assessment. A justified security migration program can therefore coexist with a decision to decline a weak quantum experiment. Four claim classes make the evidentiary distinction explicit:

1) Theoretical advantage: a proved scaling or query result under a specified model.

2) Experimental advantage: a device outperforms an identified comparator on a stated task and metric.

3) Practical advantage: the complete workflow improves a decision-relevant measure against the strongest feasible classical alternative.

4) Commercial utility: the improvement remains valuable after accounting for cost, integration, reliability, security, and opportunity cost.

Evidence supporting one claim class does not establish the next. This rule is especially important for open finance, where production claims affect regulated decisions, customer trust, and third-party dependencies.

B. Standardization is not migration

The final NIST standards establish current algorithm specifications [10]–[12], but deployment involves more than replacing a cryptographic primitive. KEM interfaces and security properties differ from traditional key exchange, while larger keys, ciphertexts, and signatures can affect fragmentation, storage, certificate chains, state machines, round trips, and latency [9], [13]. OpenPGP standardization illustrates how concrete object sizes and composite designs become protocol concerns rather than abstract algorithm choices [21]. ENISA's integration study similarly identifies effects across interfaces, protocols, operations, and hybrid designs [22].

Implementation assurance is separate from mathematical specification. Leakage studies on Kyber or ML-KEM implementations show that platform, masking, attack point, and countermeasure design affect side-channel behavior [23], [24]. A declaration that a product "supports PQC" is consequently incomplete without parameter sets, implementation provenance, conformance, side-channel evidence, interoperability results, failure behavior, and lifecycle commitments.

For an open-finance service, discovery must extend beyond transport encryption. Material uses include customer and workload authentication, API signing, consent and authorization tokens, certificate chains, secure boot, package and container signing, model provenance, firmware updates, HSMs, archival signatures, distributed ledgers where applicable, and supplier-to-supplier channels. Cryptographic inventory is therefore a graph problem, not a list of algorithm names.

C. The system boundary must include technical and institutional dependencies

The author defines quantum readiness as a system's evidenced, revisable capability to preserve digital trust, evaluate quantum-computing opportunities fairly, and govern change. This is a proposed synthesis. Assessment scope identifies the service, owner, stakeholders, components, environments, lifetimes, exclusions, and external dependencies, including those outside organizational control. Figure 1 uses seven layers: (1) information assets and protection lifetimes; (2) cryptographic mechanisms, protocols, keys, and hardware; (3) identities, trust anchors, and delegated authority; (4) applications, APIs, analytical models, and AI; (5) classical and potential quantum computing resources; (6) software delivery, provenance, and suppliers; and (7) governance, procurement, skills, and decisions.

Open standards and open source may improve inspection, portability, and shared implementation, but neither property alone proves secure composition, maintained dependencies, or operational readiness. Evidence must connect each critical function to accountable owners and change paths.

D. Missing evidence must not be averaged away

Organization-level readiness scores can conceal critical unknowns. A high average should not compensate for an unidentified system owner, unknown long-lived data, incomplete third-party discovery, untested signing path, unjustified hybrid construction, or experiment without a classical baseline. The proposed maturity scale is ordinal:

0: unknown or absent;

1: recognized but partial, manual, untested, or aspirational;

2: defined and substantially populated, with material gaps;

3: exercised in the target context with results and remediation; and

4: sustained and adaptive, with evidence of response to a material change.

Every score above zero must cite an observable artifact, test, decision, or observation attributable to a responsible person. Confidence is recorded separately. Critical conditions cap related items, and the framework does not calculate a default single composite score. The result is intended to direct discovery and action, not to manufacture precision.

IV. Proposed Framework

E. Four linked records

The four records can use existing architecture, risk, configuration, experiment, or governance tools.

System evidence graph. Links assets, identities, applications, suppliers, cryptographic uses, and dependencies within the assessment scope.

Trust-transition register. Links each security function and owner to its mechanism, lifetime, exposure, dependencies, target, tests, rollback, and residual risk.

Quantum-experiment register. Records the decision problem, data access, classical baseline, quantum method, hardware, total resources, result, claim class, and stopping decision.

Capability-and-decision register. Records maturity evidence, confidence, caps, remediation, budget, approval, exceptions, dissent, rationale, and review date.

Records carry stable identifiers, evidence links, an owner, confidence, and a verification date. Typed edges include calls, authenticates, signs, supplies, and governs. Register entries reference graph and decision identifiers. For a changed evidence item, reviewers follow relevant typed links, flag dependent entries, and record the owner and reopening condition. Graph reachability alone neither proves an impact nor authorizes a change. Section V illustrates the proposed rule; no propagation benefit has been measured.

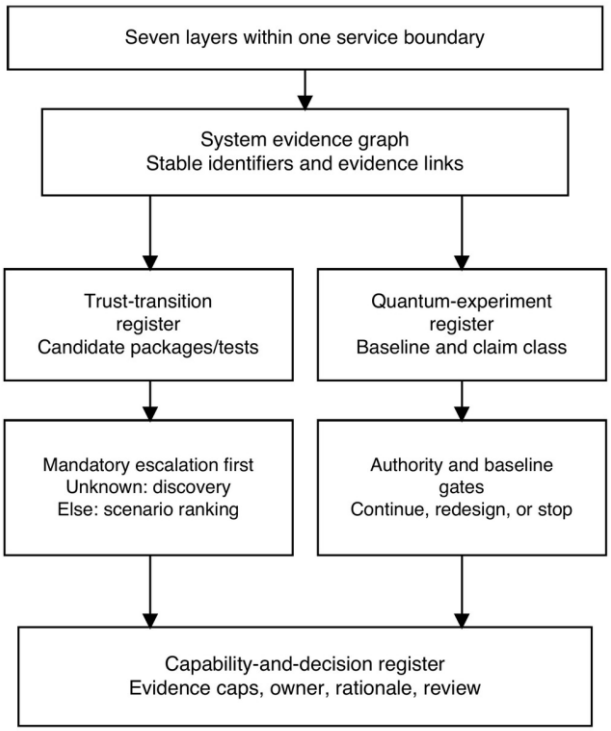


Fig. 1. Proposed four-record architecture. Shared evidence supports distinct migration and experiment gates; decisions retain links to their supporting records.

F. Migration prioritization

The unit is a candidate migration work package with an owner and security function. Criteria are protection lifetime (L), criticality (B), exposure (E), vulnerable-cryptography dependence (V), migration difficulty/lead time (M), supplier/protocol constraint (S), and agility deficit (A). For package i, criterion j, and preference scenario s, x_ij is the 0–4 input divided by four, and w_j^(s) is its declared nonnegative weight. Priority is:

$$P_i^{(s)} = \sum_{j=1}^{7} w_j^{(s)} x_{ij} \quad (1)$$

$$w_j^{(s)} \geq 0, \qquad \sum_{j=1}^{7} w_j^{(s)} = 1 \quad (2)$$

P_i^(s) expresses a preference for preparation, not a compromise probability or deployment permission. Cross-function comparison requires an explicitly agreed preference scale; it does not equate confidentiality and authenticity risks. Binding obligations, active vulnerabilities, unsupported products, expiring trust roots, and safety needs trigger mandatory escalation before ranking. Unknown scope, owner, security function, mechanism, or evidence confidence sends an item to discovery. Provisional ranking requires a documented rationale, evidence source or explicit synthetic assumption, and uncertainty range for each input. Known low confidence requires sensitivity analysis, not an automatic zero. Target, tests, and rollback may remain pending for preparation but must be resolved before deployment. Feasibility is recorded separately. Overlap among M, S, and A requires justification and criterion-removal checks.

G. Evidence-gated quantum experiments

Quantum experiments progress through seven gates:

1) Authority and safety: establish data permissions, privacy/security controls, and spending, publication, and stop authority.

2) Decision problem: specify the user, decision, constraints, instances, current method, and success metric.

3) Classical baseline: justify the strongest feasible algorithm, tuning, implementation, hardware, quality, runtime, cost, and energy accounting.

4) Quantum suitability: specify encoding, input/output, errors, scaling assumptions, and simulator or hardware limits.

5) Execution design: preregister instances, repetitions, uncertainty, resources, failures, and reproducibility requirements.

6) Fair comparison: include preparation, compilation, loading, queuing, communication, mitigation, post-processing, failure, and integration costs.

7) Decision: apply predeclared continue, redesign, or stop thresholds and assign a defensible claim class.

Hybrid workflow research shows that orchestration and provenance can improve selected runtimes without proving quantum computational advantage [25], [26]. Stopping a weak experiment is therefore evidence of governance capability, not program failure.

V. Application to Open Finance

H. Worked service trace

All architecture details, lifetimes, scores, and decisions are synthetic assumptions. The read-only bank-account aggregation service illustrates the banking subset of open

finance [1]–[3], not insurance, pensions, investments, or payment execution. A fraud-analysis model is included as a dependent consumer of authorized account-data features; no production customer data are used. Table 1 identifies the components of GRAPH1.

TABLE 1. HYPOTHETICAL SERVICE COMPONENTS AND DEPENDENCIES.

Identifiers	Role and assumed dependence
CONS1, ID1, IAM1, API1	Consent authorizes the partner workload's gateway call; the identity service authenticates it.
API1, PROV1, DATA1	Gateway retrieves account data from the provider over ECDHE-based TLS.
SIGN1, AUDIT1	Audit service signs access evidence through an RSA signature chain.
BUILD1, MODEL1	Delivery pipeline signs the fraud model's deployment artifact; signing dependencies are unknown.
SUP1	Supplier provides the cryptographic library and HSM support.

Trust-transition package WP1, owned by the API platform lead, addresses API1–PROV1 key establishment and seven-year confidentiality for DATA1. A target protocol profile still requires supplier, interoperability, payload, latency, downgrade, and rollback evidence. WP2, owned by the records lead, addresses SIGN1–AUDIT1 signing and ten-year verification, with a restricted signer and coordinated verification-profile changes across institutions. Unknown BUILD1–MODEL1 signing dependencies enter DISC1, without a priority score.

The service risk owner maintains DEC1, linking GRAPH1, both packages, and DISC1, with low confidence and an illustrative September 5, 2026 verification date. EVID1 contains assumed inventory and target-environment forward-test results but no rollback result. The transition, coexistence, and rollback maturity item is capped at level 2, despite partial level-3 evidence. DEC1 requests rollback evidence and reopens when it arrives; this does not establish organizational maturity.

EXP1 links DATA1, MODEL1, and DEC1 to quantum-assisted fraud-model optimization. Data permissions and the decision problem are assumed defined, but a justified classical baseline is absent. Gate 3 stops execution pending baseline evidence; no algorithmic feasibility or performance is established.

Consider SUP1 withdrawing library support. With identical evidence, an inventory and risk/change register can identify affected assets and migration reviews. The proposed linkage requires DEC1 to record mandatory escalation for unsupported dependencies, the unchanged rollback cap, and EXP1's separate baseline stop. The risk owner assigns mitigation review to the API and records leads before any scenario ranking; withdrawal does not reopen experiment execution. New baseline evidence would separately reopen Gate 3. Existing tools could implement these links: the contrast specifies an auditable rule, not exclusive capability, faster work, or superior decisions.

I. Reproducible priority example

These illustrative rubrics map each sequence to 0–4. L: no enduring need, up to one year, over one to three years, over three to five years, over five years. B: negligible, minor, material, major, critical consequence. E: isolated, restricted internal, institution-wide internal, private-network partner, internet-facing including authenticated APIs. V: no vulnerable dependence, ancillary, indirect, important replaceable, direct material dependence. M: routine change, single component, several components, coordinated service transition, multi-institution transition. S: no constraint, scheduled, supplier-dependent, certification-dependent, multi-party standards dependency. A: exercised replacement, tested limited replacement, defined untested path, manual partial path, no evidenced path. Unknown mandatory evidence enters discovery, not priority level 0.

The original endpoint-access illustration, before support withdrawal, uses WP1=(4,3,4,4,3,2,3) and WP2=(4,4,1,4,4,4,4) in L,B,E,V,M,S,A order. Both lifetimes exceed five years and both functions directly depend on vulnerable public-key mechanisms, giving identical L=V=4. Consequences are assumed major versus critical. E=4 describes the API endpoint; E=1 describes signer access, not public-key, artifact, or verifier exposure. Quantum forgery need not involve access to the signer. Operationally, missing verifier/artifact exposure sends WP2 to discovery. The calculations diagnose the endpoint-only proxy; they do not endorse an actionable ranking.

For WP1, assume service-level coordination (M=3), supplier dependence (S=2), and a manual partial replacement path (A=3). For WP2, assume multi-institution coordination, a multi-party profile dependency, and no evidenced replacement path (all 4). M proxies coordination/lead time, S identifies an external dependency, and A describes replacement evidence. Higher burden raises preparation priority because work must start earlier; it does not establish feasibility. Their overlap and equal numerical intervals remain modeling limitations. Operational use requires owner-approved anchors and evidence ranges.

Long-lived-data weights (.25,.10,.20,.25,.05,.05,.10) yield WP1=0.9125 and WP2=0.85. Critical-service weights (.05,.25,.10,.15,.20,.20,.05) yield 0.775 and 0.925, reversing the order. The names describe preferences, not causes: identical L and V do not distinguish these packages. Mandatory escalation overrides either ranking.

All 14 single-criterion omissions were tested with renormalization. Only omitting E reverses the first scenario (0.890625 versus 1.0). Neither order reverses across 128 independent relative ±20% weight corners per scenario, renormalized with fixed inputs. However, changing only WP2 exposure to 3 gives 0.95 in the first scenario and reverses it. Declared joint score stresses also permit both orders. These deterministic checks demonstrate sensitivity to assumptions, not operational robustness; full inputs, code, and stress ranges are retained in the evaluation protocol.

J. A starting sequence

Begin with one service, its accountable owner, protection obligations, and a cryptographic bill of materials linking identities, applications, signing paths, suppliers, and lifetimes. Queue unknowns for discovery. Select a migration package with target, interoperability, fallback, rollback, and supplier evidence; separately require claim classification and a current classical baseline for each experiment. Record confidence, exceptions, funding, and review triggers in a dated decision.

India's 2026 quantum-safe report emphasizes inventory, procurement, testing, phased migration, skills, and governance [27]. The Reserve Bank of India's CBDC concept note places quantum resistance alongside resilience, recoverability, interoperability, and vendor lock-in [28]. This exploratory

note predates final NIST standards; it is neither operational evidence nor a binding migration rule.

VI. Discussion and Evaluation Plan

Table 2 traces source findings to proposed controls. Graphs, identifiers, external-dependency analysis, and risk reviews are inherited mechanisms. The claimed increment is their explicit linkage to experiment gates and evidence-capped decisions. The inspected OpenFi abstract concerns blockchain-based financial cryptography [29]; its detailed migration and governance coverage was not assessed. Uninspected features cannot support an absence or novelty claim.

Table 2. Source-to-design trace for the four records.

Inspected basis	Design requirement	Proposed record and rule
ETSI dependency analysis [15], Secs. 6.3, 6.5	Retain internal and external dependencies	Graph links evidence and responsible owners
NIST agility [14]; ETSI migration/hybrid guidance [15], [16]	Preserve continuity through transition	Trust register links tests, rollback, and review
Readiness [17], [18]; workflow provenance [26]	Separate feasibility and workflow cost from advantage	Experiment register enforces baseline and comparison gates
ETSI priority planning [15], Sec. 6.8; business readiness [18]	Retain accountable, revisable decisions	Decision register links evidence, caps, and reopening conditions

The table supports the requirements, not the exact caps, weights, or claimed utility of joining records. Those are design propositions. The prepared evaluation protocol retains two contrasting cases, at least eight packages, four weight scenarios, and one hypothetical experiment per case. Subject to ethics, case authority, and consent, 8–12 experts would assess relevance, clarity, observability, and burden. Agreement measures, missing-data rules, and sensitivity conventions would be fixed before collection. Comparison with an inventory/risk register and generic checklist would assess traceability, missing evidence, time, and clarification burden. Section V does not complete this protocol.

The framework remains unvalidated. Limited search coverage and source access may omit closer prior work. Additive scores permit compensation after gates, overlapping criteria can amplify constraints, and cross-function comparisons depend on agreed preferences. The banking example does not establish transferability to other financial products. Two future cases and a small purposive panel could support refinement, not population-level effects or security certification. Sustained utility requires field observation.

VII. Conclusion

The four-record framework specifies how migration evidence, experiment gates, and accountable reviews can remain linked without sharing one readiness score. The synthetic banking trace shows discovery, a rollback cap, separate reopening conditions, and preference-sensitive rankings. Exposure sensitivity cautions against interpreting the rankings as calibrated risk. The framework offers inspectable decision mechanics; its practical utility remains to be evaluated.

Declarations

Funding and competing interests: This work was supported by Batoi. The author declares no competing interests beyond the disclosed affiliation and funding.

Data and ethics: No empirical dataset was collected or analyzed. Numerical inputs are constructed assumptions, not measured observations. Public sources are identified in the references. The work involved no human participants, personal data, confidential organizational data, or evaluated deployment; institutional ethics approval was not sought.

Author contributions and originality: The author was responsible for conceptualization, investigation, methodology, writing, review, and final approval. This manuscript is original, has not been published previously, and is not under consideration elsewhere.

AI Use: AI-enabled tools were used for the literature review, manuscript preparation, programming, and calculations. The author reviewed and revised the resulting material, checked the cited sources, substantive claims, and calculations, and accepts responsibility for the final manuscript.

References

[1] OECD, Shifting from Open Banking to Open Finance: Results from the 2022 OECD Survey on Data Sharing Frameworks, OECD Business and Finance Policy Papers, no. 24, 2023, doi: 10.1787/9f881c0c-en.

[2] Basel Committee on Banking Supervision, Report on Open Banking and Application Programming Interfaces, Bank for International Settlements, Nov. 2019. [Online]. Available: https://www.bis.org/bcbs/publ/d486.pdf

[3] Consultative Group on Innovation and the Digital Economy, Enabling Open Finance Through APIs, Bank for International Settlements, Dec. 2020. [Online]. Available: https://www.bis.org/publ/othp36.pdf

[4] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in Proc. 35th Annu. Symp. Foundations of Computer Science, 1994, pp. 124–134, doi: 10.1109/SFCS.1994.365700.

[5] L. K. Grover, "A fast quantum mechanical algorithm for database search," in Proc. 28th Annu. ACM Symp. Theory of Computing, 1996, pp. 212–219, doi: 10.1145/237814.237866.

[6] F. Arute et al., "Quantum supremacy using a programmable superconducting processor," Nature, vol. 574, pp. 505–510, 2019, doi: 10.1038/s41586-019-1666-5.

[7] F. Pan, K. Chen, and P. Zhang, "Solving the sampling problem of the Sycamore quantum circuits," Physical Review Letters, vol. 129, no. 9, art. 090502, 2022, doi: 10.1103/PhysRevLett.129.090502.

[8] K. Bharti et al., "Noisy intermediate-scale quantum algorithms," Reviews of Modern Physics, vol. 94, no. 1, art. 015004, 2022, doi: 10.1103/RevModPhys.94.015004.

[9] A. Banerjee, T. Reddy K., D. Schoinianakis, T. Hollebeek, and M. Ounsworth, Post-Quantum Cryptography for Engineers, RFC 9958, Jun. 2026, doi: 10.17487/RFC9958.

[10] National Institute of Standards and Technology, Module-Lattice-Based Key-Encapsulation Mechanism Standard, FIPS 203, 2024, doi: 10.6028/NIST.FIPS.203.

[11] National Institute of Standards and Technology, Module-Lattice-Based Digital Signature Standard, FIPS 204, 2024, doi: 10.6028/NIST.FIPS.204.

[12] National Institute of Standards and Technology, Stateless Hash-Based Digital Signature Standard, FIPS 205, 2024, doi: 10.6028/NIST.FIPS.205.

[13] G. Alagic et al., Recommendations for Key-Encapsulation Mechanisms, NIST SP 800-227, 2025, doi: 10.6028/NIST.SP.800-227.

[14] E. Barker et al., Considerations for Achieving Crypto Agility: Strategies and Practices, NIST CSWP 39-upd1, Dec. 2025, updated Jun. 29, 2026, doi: 10.6028/NIST.CSWP.39-upd1.

[15] European Telecommunications Standards Institute, A Repeatable Framework for Quantum-Safe Migrations, ETSI TR 104 016 V1.1.1, Oct. 2024. [Online]. Available: https://www.etsi.org/deliver/etsi_tr/104000_104099/104016/01.01.01_60/tr_104016v010101p.pdf

[16] European Telecommunications Standards Institute, Deployment Considerations for Hybrid Schemes, ETSI TR 103 966 V1.1.1, Oct. 2024. [Online]. Available:

https://www.etsi.org/deliver/etsi_tr/103900_103999/103966/01.01.01_60/tr_103966v010101p.pdf

[17] A. Purohit, M. Kaur, Z. C. Seskir, M. T. Posner, and A. Venegas-Gomez, "Building a quantum-ready ecosystem," IET Quantum Communication, vol. 5, no. 1, pp. 1–18, 2024, doi: 10.1049/qtc2.12072.

[18] OECD, Building Business Readiness for Quantum Computing: Key Barriers and Support Mechanisms, OECD Digital Economy Papers, no. 383, 2026, doi: 10.1787/ee847e5f-en.

[19] A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design science in information systems research," MIS Quarterly, vol. 28, no. 1, pp. 75–105, 2004, doi: 10.2307/25148625.

[20] K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, "A design science research methodology for information systems research," Journal of Management Information Systems, vol. 24, no. 3, pp. 45–77, 2007, doi: 10.2753/MIS0742-1222240302.

[21] S. Kousidis, J. Roth, F. Strenzke, and A. Wussler, Post-Quantum Cryptography in OpenPGP, RFC 9980, 2026, doi: 10.17487/RFC9980.

[22] D. J. Bernstein, A. Huelsing, and T. Lange, Post-Quantum Cryptography - Integration Study, European Union Agency for Cybersecurity, 2022, doi: 10.2824/151162.

[23] Z. Huang, H. Wang, B. Cao, D. He, and J. Wang, "A comprehensive side-channel leakage assessment of CRYSTALS-Kyber in IIoT," Internet of Things, vol. 27, art. 101331, 2024, doi: 10.1016/j.iot.2024.101331.

[24] E. Camacho-Ruiz, P. Navarro-Torrero, and A. Cabrera Aldaya, "A framework for designing high-order side-channel protected hardware implementations of ML-KEM," IACR Transactions on Cryptographic Hardware and Embedded Systems, vol. 2026, no. 2, pp. 272–295, 2026, doi: 10.46586/tches.v2026.i2.272-295.

[25] A. A. Khan, A. Ahmad, M. Waseem, P. Liang, M. Fahmideh, T. Mikkonen, and P. Abrahamsson, "Software architecture for quantum computing systems - A systematic review," Journal of Systems and Software, vol. 201, art. 111682, 2023, doi: 10.1016/j.jss.2023.111682.

[26] B. Weder, J. Barzen, M. Beisel, and F. Leymann, "Provenance-preserving analysis and rewrite of quantum workflows for hybrid quantum algorithms," SN Computer Science, vol. 4, art. 233, 2023, doi: 10.1007/s42979-022-01625-9.

[27] Department of Science and Technology, Government of India, Quantum-Safe Ecosystem in India, National Quantum Mission, May 2026. [Online]. Available: https://dst.gov.in/sites/default/files/Quantum-Safe-Ecosystem-in-India.pdf

[28] Reserve Bank of India, Concept Note on Central Bank Digital Currency, Oct. 7, 2022. [Online]. Available: https://www.rbi.org.in/Scripts/PublicationReportDetails.aspx?ID=1218

[29] S. Banaeian Far, M. R. Chalak Qazani, and A. Imani Rad, "Exploring blockchain-based open finance: The role of advanced cryptographic tools, state-of-the-art achievements, and background challenges," Peer-to-Peer Networking and Applications, vol. 18, art. 157, 2025, doi: 10.1007/s12083-025-01979-w.