

Post-Quantum Cryptography: Evaluating Transition Strategies for Legacy Network Infrastructure

Nidhi Srivastava

Nidhiaransh@gmail.com

<https://orcid.org/0009-0002-5384-7873>

Abstract—Large-scale quantum computers would break the public-key algorithms that secure most of today's networks, including RSA, Diffie-Hellman and elliptic-curve cryptography. Standards bodies have responded with post-quantum algorithms, but the harder problem is migration: legacy network infrastructure contains long-lived devices, constrained hardware, embedded certificates and protocols that were never designed for large keys or signatures. This paper evaluates transition strategies for such environments. We summarise the quantum threat and the main families of post-quantum algorithms, analyse the operational constraints of legacy systems, and compare five strategies: direct replacement, hybrid classical and post-quantum schemes, crypto-agility, gateway-based overlays, and quantum-assisted channels. Using a digital-signature workflow and the key-establishment flow of the Messaging Layer Security protocol as worked examples, we show where algorithm changes touch protocol design. A simple throughput model for mixed post-quantum and quantum-channel traffic illustrates how capacity degrades as more transactions use slower mechanisms. We conclude with a phased roadmap, governance guidance and open research questions.

Keywords — Post-Quantum Cryptography, Legacy Infrastructure, Hybrid Key Exchange, Crypto-Agility, ML-KEM, ML-DSA, Migration, Quantum Threat.

I. Introduction

Public-key cryptography underpins nearly every secure network interaction. It authenticates servers and software updates, establishes the session keys used by TLS, SSH and VPNs, and anchors the certificate hierarchies of the public key infrastructure. The security of the most widely deployed schemes, RSA [1], finite-field Diffie-Hellman [2] and elliptic-curve variants, depends on the difficulty of factoring large integers or computing discrete logarithms. In 1994 Shor showed that a sufficiently large quantum computer could solve both problems in polynomial time [3]. Grover's search algorithm, in contrast, gives only a quadratic speed-up against symmetric ciphers and hash functions [4], which can be countered by increasing key and output lengths.

No cryptographically relevant quantum computer exists today, and expert estimates of when one might appear vary widely. Nevertheless, the threat is already operational because of harvest-now-decrypt-later attacks, in which an adversary records encrypted traffic today and decrypts it once quantum capability becomes available. Data that must remain confidential for decades, such as health records, state secrets, intellectual property and infrastructure designs, is therefore at risk now. Digital signatures face a different but related risk: a forged signature can only be exploited once a quantum computer exists, yet systems that rely on long-lived signatures, such as firmware roots of trust, must be protected before that moment.

The cryptographic community has responded through a multi-year standardisation effort led by the U.S. National Institute of Standards and Technology (NIST), which in August 2024 published its first post-quantum standards: FIPS 203 for the ML-KEM key-encapsulation mechanism, FIPS 204 for the ML-DSA signature scheme and FIPS 205 for the hash-based SLH-DSA signature scheme [5]–[7]. A further signature standard based on Falcon, to be named FN-DSA, and a code-based key-encapsulation mechanism, HQC, selected in 2025 as a backup to ML-KEM, were still being finalised at the time of writing. NIST has also published a transition plan proposing that quantum-vulnerable algorithms be deprecated and ultimately disallowed in the coming decade [8].

Having standards is not the same as having a migrated network. Legacy infrastructure, meaning the routers, firewalls, VPN concentrators, load balancers, industrial controllers, payment terminals, smart

10.48047/jocaaa.2025.34.08.19

cards and embedded devices that form the backbone of enterprise and critical-infrastructure networks, was designed around small keys, compact signatures and long service lives. Post-quantum algorithms generally require larger keys and signatures, more memory and different performance trade-offs, and many legacy components cannot be upgraded in place.

This paper evaluates how organisations can navigate that gap. Its contributions are: (i) a concise taxonomy of post-quantum algorithm families and their suitability for legacy contexts; (ii) an analysis of the technical and organisational constraints of legacy networks; (iii) a comparison of five transition strategies with explicit trade-offs; (iv) worked examples using a signature workflow, the MLS key-establishment protocol and a throughput model; and (v) a phased migration roadmap with governance guidance.

II. Background

A. The Quantum Threat to Public-Key Cryptography

Shor's algorithm threatens every scheme whose security rests on integer factorisation or on discrete logarithms in finite fields or elliptic-curve groups [3]. This includes key exchange, public-key encryption and digital signatures. Symmetric primitives are affected less. Grover's algorithm reduces the effective security of a k -bit symmetric key to roughly $k/2$ bits [4], so AES-256 and SHA-384 or SHA-512 are generally regarded as adequate for the quantum era, whereas 128-bit keys offer a smaller margin. The practical priority for migration is therefore public-key mechanisms.

B. Timing Risk: The Migration Inequality

Mosca formalised the timing problem with a simple inequality [9]. Let X be the number of years that data must remain secure, Y the number of years needed to migrate the relevant systems to quantum-safe protection, and Z the number of years until a cryptographically relevant quantum computer exists. If

$$X + Y > Z$$

then the data will be exposed before protection is complete. The inequality has two implications for legacy networks. First, Y is large because inventories are incomplete and equipment lifecycles are long, so even moderate estimates of Z justify starting early. Second, X differs sharply across assets, which allows migration to be prioritised by data lifetime and sensitivity rather than attempted everywhere at once.

C. Standards Landscape

The first NIST standards cover key establishment and digital signatures. ML-KEM, derived from CRYSTALS-Kyber [10], is the primary key-encapsulation mechanism and is based on module lattices. ML-DSA, derived from CRYSTALS-Dilithium [11], is the primary signature scheme. SLH-DSA, derived from SPHINCS+ [12], relies only on hash functions and serves as a conservative alternative with larger signatures. Complementary guidance on stateful hash-based signatures, which are suited to firmware signing where state can be managed, is provided in NIST SP 800-208 [13]. Protocol-level standardisation in the IETF is progressing in parallel, including hybrid key exchange for TLS 1.3 [14] and the use of hybrid public-key encryption in the Messaging Layer Security protocol [15], [16].

III. Post-Quantum Algorithm Families

Post-quantum cryptography refers to classical algorithms believed to resist both classical and quantum attack [17]. Fig. 1 shows the principal families, and Table I summarises their properties for deployment in legacy contexts.

10.48047/jocaaa.2025.34.08.19




Fig. 1. Families of post-quantum cryptographic algorithms: lattice-based, hash-based, isogeny-based, code-based, multivariate and MPC-in-the-head.

A. Lattice-Based

Lattice schemes base security on problems such as Learning With Errors and its structured variants. They offer a favourable balance of speed and size and underpin ML-KEM, ML-DSA and Falcon. Their implementations use polynomial arithmetic that is efficient on general-purpose processors, which makes them the default choice for most migration plans. Their keys and signatures remain larger than those of elliptic-curve cryptography by one to two orders of magnitude.

B. Hash-Based

Hash-based signatures derive security solely from the properties of cryptographic hash functions, which gives them the most conservative security assumptions. SLH-DSA is stateless but produces large signatures and is slower to sign. Stateful schemes such as LMS and XMSS have smaller signatures but require careful state management to avoid catastrophic key reuse, which restricts them to controlled settings such as firmware signing [13].

C. Code-Based, Multivariate, Isogeny-Based and MPC-in-the-Head

Code-based schemes rely on the hardness of decoding error-correcting codes. The classic McEliece system has very large public keys but a long security record, and HQC, a code-based key-encapsulation mechanism, has been selected as a diversification option alongside lattice-based ML-KEM. Multivariate schemes use systems of quadratic equations and often yield short signatures, yet several candidates have been broken, notably the Rainbow scheme [18]. Isogeny-based schemes offer very compact keys, but the SIDH protocol was broken by a classical attack in 2022 [19], illustrating the risk of adopting young assumptions. Signature designs that apply the MPC-in-the-head paradigm derive security from multi-party computation protocols simulated by the signer and are an active area of research. The episodes involving Rainbow and SIDH support a central lesson for migration: algorithm choices must be revisable, which motivates crypto-agility.

TABLE I. POST-QUANTUM FAMILIES AND SUITABILITY FOR LEGACY CONTEXTS

Family	Standardised or leading examples	Strengths	Concerns
Lattice	ML-KEM, ML-DSA, Falcon (FN-DSA)	Fast, balanced sizes, broad protocol support	Larger keys than elliptic curves; newer assumptions
Hash-based	SLH-DSA, LMS, XMSS	Minimal assumptions, conservative	Large signatures; state management for LMS and XMSS
Code-based	HQC, Classic McEliece	Long cryptanalytic history, diversity from lattices	Large keys or ciphertexts
Multivariate	Rainbow (broken), newer designs	Short signatures	History of practical breaks

Isogeny-based	SIDH (broken), newer designs	Very compact keys	Young assumptions; slower
MPC-in-the-head	Research signature candidates	Security from simple primitives	Larger signatures; maturing

IV. Constraints of Legacy Network Infrastructure

A. Size and Bandwidth Overheads

Table II compares representative sizes of classical and post-quantum primitives. An ML-KEM-768 public key is 1,184 bytes and its ciphertext 1,088 bytes, compared with 32 bytes for an X25519 public value. An ML-DSA-65 signature is 3,309 bytes against 64 bytes for Ed25519. These differences matter in protocols with fixed-size fields, small maximum transmission units, or limited handshake buffers. A TLS handshake that carries a certificate chain with several post-quantum signatures and public keys may span many packets, increase latency and trigger fragmentation problems in middleboxes that mishandle large messages. Performance studies of post-quantum authentication in TLS 1.3 have shown that handshake size and the choice of signature algorithm dominate the overhead [20].

TABLE II. REPRESENTATIVE SIZES IN BYTES OF CLASSICAL AND POST-QUANTUM PRIMITIVES

Algorithm	Public key	Signature or ciphertext	Type
X25519	32	32	Classical key exchange
RSA-2048	256	256	Classical signature
Ed25519	32	64	Classical signature
ML-KEM-768	1,184	1,088	Post-quantum KEM
ML-DSA-65	1,952	3,309	Post-quantum signature
FN-DSA-512 (Falcon-512)	897	about 666	Post-quantum signature
SLH-DSA-128s	32	7,856	Post-quantum signature
SLH-DSA-128f	32	17,088	Post-quantum signature

B. Computational and Memory Limits

Lattice operations are efficient on modern servers, yet constrained devices such as smart cards, sensors and microcontrollers may lack the memory for large keys and the stack space for signature generation. Hardware security modules and network accelerators often implement fixed cryptographic engines in firmware or silicon, and cannot add new algorithms without replacement or vendor updates. Side-channel resistance must also be re-established for new algorithms, because implementation attacks target the specific arithmetic of each scheme.

C. Protocol and Format Rigidity

Protocols and file formats frequently hard-code algorithm identifiers, field sizes and certificate profiles. X.509 parsers, code-signing pipelines, smart-card applications and industrial protocols may reject unfamiliar object identifiers or oversized fields. In deployed systems the cryptographic algorithm is often entangled with application logic rather than abstracted behind an interface, so a change that is conceptually simple requires code modification and re-certification.

D. Lifecycle and Operational Realities

Operational technology, medical devices, telecommunications equipment and vehicles can remain in service for fifteen to thirty years, and firmware update channels may be absent or require physical access. Certification regimes can make changes slow and expensive. Moreover, organisations rarely have a complete inventory of where cryptography is used, including embedded libraries and third-party dependencies, so the first practical obstacle is discovery rather than algorithm selection.

E. Interoperability and Ecosystem Dependencies

Networks do not migrate in isolation. Partners, customers, cloud providers and public certificate authorities move on their own schedules, and a connection is only as quantum-safe as the weaker side permits. Negotiation mechanisms allow a client and server to select the strongest mutually supported option, but they also create downgrade risk if a classical-only fallback is silently accepted.

10.48047/jocaaa.2025.34.08.19

Organisations should therefore track the readiness of external dependencies, define policy for when fallback is allowed, and log negotiated algorithms so that residual classical sessions can be measured and reduced.

V. Transition Strategies

We compare five strategies that can be combined across an estate. Table III summarises their strengths and weaknesses.

TABLE III. COMPARISON OF TRANSITION STRATEGIES

Strategy	Main benefit	Main drawback	Best suited to
Direct replacement	Simplicity, no dual-algorithm overhead	No fallback if the algorithm is weakened	Controlled new systems
Hybrid	Security holds if either component survives	Larger messages, more complexity	External endpoints, VPNs, long-lived data
Crypto-agility	Makes any change cheap and reversible	Requires architecture and process investment	All systems; precondition for others
Gateway overlay	Protects unmodifiable legacy devices	Exposure behind the gateway; single point of dependency	OT, embedded and end-of-life systems
Quantum-assisted channels	Security from physics, not computational assumptions	Specialised hardware, range and rate limits	Few high-value links

A. Direct Replacement

Direct replacement substitutes a post-quantum algorithm for a classical one. It is simple to reason about and eliminates the complexity of running two schemes, but it forgoes protection if the new algorithm proves weak, as the Rainbow and SIDH results demonstrate [18], [19]. It is most defensible where a mature standardised algorithm exists, the environment is controlled, and interoperability with legacy peers is not required.

B. Hybrid Schemes

Hybrid schemes combine a classical algorithm with a post-quantum one so that security holds as long as at least one component remains unbroken. For key establishment, the two shared secrets are combined through a key-derivation function, as in hybrid key exchange for TLS 1.3 [14], and the combiner design has been analysed formally [21]. For authentication, composite or dual signatures can bind both algorithms. Hybrids are the prevailing recommendation during the transition because they hedge against immaturity of new algorithms and against implementation flaws, and they preserve compliance with regimes that still require classical algorithms. Their costs are larger messages, additional computation and greater complexity in negotiation and certificate handling.

C. Crypto-Agility

Crypto-agility is the architectural capability to replace cryptographic algorithms, parameters and libraries without redesigning the system. It involves abstracting cryptography behind well-defined interfaces, avoiding hard-coded identifiers, supporting algorithm negotiation, maintaining a cryptographic bill of materials, and automating certificate and key rotation. Agility is not a single migration step but a precondition for all strategies, because it makes both the initial move and any subsequent change inexpensive. For legacy systems, the realistic objective is often to introduce agility at the boundaries, in gateways and libraries, even if the core cannot be modified.

D. Gateway and Overlay Approaches

Where the legacy endpoint cannot be changed, protection can be added around it. A post-quantum-capable VPN, TLS-terminating proxy or encrypting gateway can secure the traffic between sites or between a client and the edge, leaving the legacy device untouched behind the gateway. Pre-shared symmetric keys can augment existing VPN tunnels, as a stop-gap, to provide quantum resistance without new public-key algorithms. The disadvantage is that protection stops at the gateway, so the segment behind it remains exposed, and the gateway becomes a critical dependency. This approach fits the zero-trust principle of enforcing security at defined points while systems are modernised.

E. Quantum-Assisted Channels

A different path uses quantum mechanics defensively. Quantum key distribution (QKD) establishes keys whose secrecy is guaranteed by physical principles rather than computational assumptions [22]. It requires dedicated hardware and optical links, has distance and rate limitations, and still needs authentication, typically through pre-shared keys or post-quantum signatures. It can complement post-quantum cryptography for a small number of high-value links but does not scale to general-purpose networks. We examine its throughput implications in Section VI.

F. Example 1: The Digital Signature Workflow

Signatures illustrate why migration is more than a library swap. Fig. 2 shows the standard workflow: the sender hashes the plain text, encrypts the hash with the sender's private key to produce the digital signature, appends it to the message and transmits the signed message. The receiver detaches the signature, decrypts it with the sender's public key, independently hashes the received text and accepts the message if the two hashes match.




Fig. 2. Digital signature workflow: hash and sign with the private key, append, transmit, then detach, decrypt with the public key and compare hashes.

The description uses the textbook formulation of signing as private-key encryption of a hash, which maps directly onto RSA. Post-quantum schemes such as ML-DSA and SLH-DSA follow the same logical pattern of hashing, signing with a private key and verifying with a public key, but internally use different constructions and are not literally encryption operations. The impact on legacy systems arises from three details. The signature field grows from tens or hundreds of bytes to thousands, so formats and buffers must accommodate it. The hash function and its output length may change, and signing may be randomised or require large internal state. Finally, the verifier must know which algorithm was used, so algorithm identifiers and certificate profiles must be updated. Firmware and software signing are the highest-priority use cases because devices verify signatures for years after they are issued.

G. Example 2: Key Establishment in MLS

Fig. 3 illustrates the Messaging Layer Security protocol, in which a host A adds a client B to a group. Each participant holds a credential bound to a signature key, verified by an authentication service, and publishes a KeyPackage through a directory. A generates random secrets, derives a node secret and a joiner secret using HKDF, and encrypts the path secret, joiner secret and any pre-shared key to B's public key. The encrypted welcome travels through the delivery service, and B decrypts it with its private key to join the group's key schedule.




Fig. 3. MLS key establishment between host A and client B using an authentication service, a directory of KeyPackages and a delivery service.

The figure shows two distinct uses of public-key cryptography. Credentials and signatures authenticate identity, while public-key encryption of secrets under B's key provides confidentiality for key establishment. In a post-quantum migration the encryption step would be realised with a key-encapsulation mechanism, which the MLS standard accesses through hybrid public-key encryption [15], [16]. The symmetric HKDF derivations are comparatively unaffected, given adequate output lengths. The example demonstrates two principles. First, each KeyPackage must now carry a larger public key and signature, so the directory and delivery service must handle larger objects. Second, because MLS negotiates ciphersuites, a hybrid or post-quantum ciphersuite can be introduced alongside classical ones, which is crypto-agility in practice. Group messaging also confronts harvest-now-decrypt-later directly, since recorded group traffic could be decrypted retroactively if the initial key establishment relies on classical algorithms alone.

H. Choosing Among the Strategies

The strategies are selected per asset rather than per organisation. A useful decision sequence begins with data lifetime. Assets that protect long-lived secrets and are reachable by an adversary who can record traffic should receive hybrid key exchange first. Next, ask whether the endpoint can be modified. If it can, update its libraries and enable agility; if it cannot, place a gateway in front of it and schedule replacement. Then consider signatures: anything that is verified for many years, such as firmware, boot loaders and root certificates, should move early to a quantum-safe signature, whereas short-lived signatures such as ephemeral session authentication can follow once protocol support matures. Quantum-assisted channels are considered last and only for links that justify the cost. Recording the reasoning for each decision supports audit and later revision.

VI. Performance Modelling of Mixed Protection Mechanisms

Transition strategies change capacity as well as security. We illustrate this with a simple model of a service that processes transactions protected by either post-quantum mechanisms or a slower quantum-assisted channel. Let κ_Q denote the fraction of transactions that use the quantum-assisted channel, T_{PQC} the throughput when all transactions use post-quantum cryptography, and T_Q the throughput when all use the quantum channel. The aggregate throughput in transactions per second is then the weighted average

10.48047/jocaaa.2025.34.08.19

$$T(kappa_Q) = (1 - kappa_Q) \times T_{PQC} + kappa_Q \times T_Q$$

Fig. 4 plots this relationship for T_{PQC} of 1,000 transactions per second and three values of T_Q . In the central scenario T_Q is 200, and the plotted sample points are consistent with the model: 920 transactions per second when ten percent of transactions use the quantum channel, 600 at fifty percent and 200 when all do. The two other scenarios, with T_Q of 10 and 500, bound an uncertainty band that reflects how little is known about the sustained rates of quantum-assisted systems. The horizontal dashed line marks the quantum-only throughput of the central scenario, and the post-quantum-only case lies at the left edge of the graph where $kappa_Q$ is zero.




Fig. 4. Throughput versus fraction of quantum transactions for three quantum-channel scenarios, with the uncertainty band shaded.

A. Interpretation

- **Throughput falls linearly as the quantum share grows.** Every additional ten percentage points of quantum traffic reduces capacity by one tenth of the gap between the two mechanisms, which is eighty transactions per second in the central scenario.
- **The uncertainty is largest at high quantum share.** At ten percent the three scenarios differ by only about ninety transactions per second, but at one hundred percent they range from 10 to 500, a fifty-fold difference. Small deployments of quantum channels are therefore low risk for capacity, whereas large ones depend heavily on an uncertain parameter.
- **Selective use is the practical conclusion.** Quantum-assisted channels should be reserved for a narrow set of high-value links, while post-quantum cryptography protects the bulk of traffic.

B. Limits of the Model

The linear form assumes that the two mechanisms operate on independent resources, so that capacity can be divided in proportion to traffic share. If both share a single processing resource, with each transaction consuming a time cost, the aggregate rate is the harmonic combination, $1 / ((1 - kappa_Q) / T_{PQC} + kappa_Q / T_Q)$, which falls faster than the linear model for mixed workloads. The model also ignores latency, handshake amortisation through session reuse, queueing effects and the cost of key management. It should be understood as a first-order planning aid that shows the direction and relative size of the capacity effect, and it should be replaced by measurements from the organisation's own systems before capacity decisions are made. The same weighted-average logic applies to hybrid and post-quantum handshakes, where larger messages and extra computation reduce the connection rate of gateways and servers.

10.48047/jocaaa.2025.34.08.19

A practical reading of the model is that planners should estimate the two end-point throughputs on their own platforms, choose the share of traffic expected to use each mechanism, and test the sensitivity of the result to the less certain parameter. If the lower bound of the uncertainty band still meets service-level objectives at the intended share, the design is robust; if it does not, the share should be reduced or capacity added.

VII. A Phased Migration Roadmap

Combining the strategies above, we propose a four-phase roadmap suitable for legacy-heavy environments. Table IV summarises the phases and indicative outputs.

TABLE IV. PHASES OF THE MIGRATION ROADMAP

Phase	Focus	Indicative outputs
1 Discover and prioritise	Inventory and risk ranking	Cryptographic bill of materials; assets ranked by $X + Y$ versus Z
2 Build agility and pilot	Abstraction, vendor requirements, lab testing	Updated libraries; pilot results on sizes, latency and compatibility
3 Deploy hybrid protection	External links, gateways, signing roots	Hybrid key exchange live; legacy systems behind gateways; PQ firmware signing
4 Retire and optimise	Remove classical-only configurations	Decommissioned devices; policy aligned with regulatory timelines

Phase 1: Discover and Prioritise

Build a cryptographic inventory, ideally a cryptographic bill of materials, covering protocols, libraries, certificates, hardware modules and third-party dependencies. Classify each asset by data lifetime X , the estimated migration time Y and exposure to harvest-now-decrypt-later attack. Apply the inequality in Section II to rank assets, giving priority to long-lived confidential data, root-of-trust signatures, and externally exposed links.

Phase 2: Build Agility and Pilot

Introduce abstraction layers, update libraries and procurement requirements, and require post-quantum support from vendors. Establish a laboratory to test hybrid handshakes, measure packet sizes, latency and CPU load, and detect middlebox failures. Pilot on non-critical services and publish the results to inform wider rollout.

Phase 3: Deploy Hybrid Protection

Enable hybrid key exchange on external-facing endpoints, VPNs and inter-site links, starting where confidentiality lifetimes are longest. Place gateways in front of legacy systems that cannot be upgraded. Begin issuing post-quantum or composite certificates from internal certificate authorities for new systems, and move firmware-signing roots to quantum-safe signatures, using hash-based schemes where state can be managed securely.

Phase 4: Retire and Optimise

Replace or isolate end-of-life devices, retire classical-only algorithms in line with regulatory timelines, and consider moving from hybrid to post-quantum-only configurations when confidence in the new algorithms and the regulatory position justify it. Continue monitoring cryptanalytic progress and maintain the agility to switch algorithms again.

VIII. Governance, Risk and Compliance

Technical migration succeeds only with organisational support. Executive sponsorship, a named owner for cryptographic risk and a funded programme are prerequisites, since the work spans networking, application, procurement and security teams. Procurement policy should require crypto-agility and a vendor roadmap for post-quantum support, and contracts should address firmware updates over the device lifetime. Risk registers should record the harvest-now-decrypt-later exposure explicitly, and

10.48047/jocaaa.2025.34.08.19

exceptions for systems that cannot migrate should carry compensating controls and expiry dates. Compliance regimes, including regulations that mandate validated cryptographic modules, affect timing, because new algorithms may need module validation before they can be used in regulated settings. Finally, testing and incident-response plans should anticipate cryptographic failure, including the discovery of a break in a newly adopted algorithm.

IX. Limitations

This paper is a conceptual and analytical evaluation and does not report new empirical measurements of post-quantum performance on legacy hardware. Sizes quoted in Table II are taken from the algorithm specifications and may vary with parameter sets and encodings. The throughput model is a deliberately simple first-order illustration with assumed parameters, and the quantum-assisted channel rates are scenario values rather than measured figures. Standards continue to evolve: the status of FN-DSA and HQC, regulatory deadlines and protocol drafts may change after this writing, and readers should confirm current requirements. Finally, the roadmap is generic and must be adapted to each organisation's regulatory, technical and risk context.

X. Future Research Directions

Several directions deserve attention. Empirical benchmarking of hybrid and post-quantum handshakes across real legacy middleboxes, constrained devices and hardware security modules would replace assumptions with data. Compact signature designs and certificate-compression techniques could reduce handshake overhead. Automated discovery tools for cryptographic usage in binaries and firmware would shorten the inventory phase. Formal analysis of hybrid combiners and composite certificates under realistic attacker models remains active. Finally, better models of the cost and capacity of quantum-assisted networks, validated by measurement, would improve decisions about where such links are justified.

XI. Conclusion

The quantum threat to public-key cryptography is a migration problem as much as a mathematical one. Standards for post-quantum key establishment and signatures now exist, but legacy network infrastructure, with its long lifetimes, rigid protocols and constrained hardware, cannot adopt them by simple substitution. This paper compared five transition strategies and argued that they are complementary. Hybrid schemes provide a safety margin while the new algorithms mature, gateways protect systems that cannot change, quantum-assisted channels suit a narrow set of high-value links, and crypto-agility is the foundation that makes each of the others affordable and reversible.

The worked examples showed where migration touches real designs: signature fields and certificate formats grow, key-establishment messages in protocols such as MLS must carry larger objects and negotiate new ciphersuites, and capacity planning must account for slower mechanisms. The phased roadmap turns these observations into an order of work that begins with inventory and prioritisation by data lifetime, proceeds through agility and pilots, deploys hybrid protection where exposure is greatest, and ends with retirement of classical-only configurations.

Because harvest-now-decrypt-later attacks are possible today, the date at which a quantum computer arrives is not the relevant deadline for the most sensitive data. Organisations that begin the inventory now, demand agility from their suppliers and deploy hybrid protection in priority order will be best placed to complete the transition before it becomes urgent.

References

1: Pawan Kumar Singh, Scaling Gate-All-Around (GAA) Transistor Architectures for Sub-2nm AI Accelerators Using Atomically Thin 2D Materials, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper), 1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5713>

- 2: **Pawan Kumar Singh**, Energy-Efficient Task Scheduling in Green Cloud Computing Using Neuromorphic Spiking Neural Networks, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5712>
- 3: **A. Mahendra Vardhan and S. Sridhar**, "Determining False Positive Analysis of Software Vulnerabilities with Predefined Scan Rules using Random Forest Classifier and Decision Tree Technique," 2022 4th International Conference on Advances in Computing, Communication Control and Networking (ICAC3N), Greater Noida, India, 2022, pp. 622-625, DOI: <https://doi.org/10.1109/ICAC3N56670.2022.10074458>
- 4: **Ronakkumar shah** , Cognitive Process Intelligence: Integrating Generative AI and Automated Workflow Discovery, Vol. 33 No. 08 (2024): JOCAAA, Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5676>
- 5: **Ronakkumar shah**, research Trends in Smart Workflow Automation: A Bibliometric Analysis of Emerging Themes, Home / Archives / Vol. 33 No. 05 (2024): JOCAAA, Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5675>
- 6: **Dip Bharatbhai Pate**, Dynamic vs. Static visualizations: Understanding their use cases in big data analysis, International Journal of Science and Research Archive, 2022, 05(02), 392-395., DOI <https://doi.org/10.30574/ijrsra.2022.5.2.0077>
- 7: **Dip Bharatbhai Patel**, Building scalable business intelligence systems in the cloud: A technical approach, International Journal of Science and Research Archive, 2021, 02(01), 212-215., Received on 03 February 2021; revised on 17 April 2021; accepted on 19 April 2021, DOI: <https://doi.org/10.30574/ijrsra.2021.2.1.0047>
- 8: **Dip Bharatbhai Patel**, Understanding the Impact of Text Analytics on Social Media Sentiment Analysis, Vol. 5 No. 5 (2019) ,September-October, DOI: <https://doi.org/10.32628/CSEIT1936413>
- 9: **Dip Bharatbhai Patel**, Real-Time Data Streaming in BI Dashboards: Opportunities and Challenges, Archives / Vol. 5 No. 4 (2019) : July-August, DOI: <https://doi.org/10.32628/CSEIT192142>
- 10: **Dip Bharatbhai Patel**, the Role of Data Visualization in Modern Intelligence, **Archives / Vol. 2 No. 2 / Research Article, Published, 2020-11-28 DOI: <https://doi.org/10.32996/jcsts.2020.2.2.4>**
- 11: **A. Srivastava**, "Securing PHI in cloud-based healthcare systems: Challenges, frameworks, and future directions," J. Comput. Anal. Appl., vol. 33, no. 2, Aug. 2024. [Online]. Available: <https://www.eudoxuspress.com/index.php/pub/article/view/4349/3190>
- 12: **A. Srivastava**, "AI-assisted cloud migration of healthcare claims data from legacy systems: A metadata-driven framework," Int. J. Commun. Networks Inf. Secur., vol. 14, no. 3, Nov. 2022. [Online]. Available: <https://ijcnis.org/index.php/ijcnis/article/view/8682>
- 13: **A. Srivastava**, "Enhancing provider and claims data accuracy using artificial intelligence on cloud-native data platforms," J. Comput. Anal. Appl., vol. 31, no. 4, Nov. 2023. [Online]. Available: <https://eudoxuspress.com/index.php/pub/article/view/4347/3189>

10.48047/jocaaa.2025.34.08.19

14: **A. Srivastava**, "A unified framework for secure cloud data management: Integrating governance, metadata intelligence, and privacy controls," J. Inf. Syst. Eng. Manage., vol. 9, no. 2, Apr. 2024. [Online]. Available: https://www.jisem-journal.com/download/45_A_Unified_Framework.pdf

15: **A. Srivastava**, "Optimizing large-scale data migration for cloud-native architectures," J. Comput. Anal. Appl., vol. 31, no. 2, pp. 652–662, May 2023. [Online]. Available: <https://eudoxuspress.com/index.php/pub/article/view/4603>

16: **Abdul Aziz**. "Intelligent Health Analytics: Mastering AI in Healthcare Data". Acta Scientiae, vol. 25, no. 4, Aug. 2024, pp. 544-55, <https://www.periodicos.ulbra.org/index.php/acta/article/view/687>. <https://doi.org/10.5281/zenodo.21699886>

17: Event-Aware Multi-Layer Storage Risk Forecasting for Oracle Database Estates Using HAPF, Raghu Gollapudi, International Journal of Computational and Experimental Science and Engineering, 10(4), 2024 ,[10.22399/ijcesen.5183](https://www.ijcesen.com/index.php/ijcesen.5183)

18: Backup Integrity and Recovery Readiness Assessment for High-Availability Databases Raghu Gollapudi, Computer Fraud & Security, Vol. 2023, Issue 8, pp. 35-48, [10.52710/cfs.1031](https://www.cfsjournal.com/index.php/cfs.1031)

19: Risk-Controlled Near-Zero-Downtime Oracle Database Migration Using GoldenGate Raghu Gollapudi, International Journal of Computational and Experimental Science and Engineering, 8(3), 2022 ,[10.22399/ijcesen.5382](https://www.ijcesen.com/index.php/ijcesen.5382)

20: An Automated Risk Scoring Framework for SQL Execution Plan Analysis and Performance
56: Manoj Kumar Kagitha, Automated Security Compliance in Ci/Cd: A Natural Language Processing (NLP) Approach to Detecting Vulnerabilities in Infrastructure-As-Code (Terraform/Bicep), Publication year 2021, Volume 2021, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/970>

20: Manoj Kumar Kagitha, Hybrid Quantum-Classical Resource Scheduling: A Proposed Architecture for Next-Generation Hyperscale Data Centers, Publication year 2024, Volume 2024, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/969>

22: Manoj Kumar Kagitha, Self-Healing Cloud Infrastructure: A Hybrid Reinforcement Learning Framework for Predicting Microservice Failures in Azure Kubernetes Service (AKS), Publication year 2023, Volume 2023, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/971>

23: Manoj Kumar Kagitha, LOAD BALANCING TECHNIQUES IN DISTRIBUTED CLOUD ENVIRONMENTS, (IJESMR - 2017), D.O.I:10.53555/jcr.v6:i1.13380, <https://jcreview.com/archives/volume-6/issue-1/13380>

24: Manoj Kumar Kagitha, AI-DRIVEN OBSERVABILITY FOR HYPERSCALE COLOCATION DATA CENTERS: A CASE STUDY OF LOG ANALYTICS AND ANOMALY DETECTION PIPELINES, Archives / Vol. 27 No. 7 (2019): JOCAAA, <https://www.eudoxuspress.com/index.php/pub/article/view/5053>

25: Manoj Kumar Kagitha, GREENOPS IN THE CLOUD: AN AI-DRIVEN TELEMETRY ANALYSIS MODEL FOR REDUCING CARBON FOOTPRINT IN HIGH-AVAILABILITY CLUSTERS, Vol. 28 No. 6 (2020): JOCAAA, <https://www.eudoxuspress.com/index.php/pub/article/view/5054>

10.48047/jocaaa.2025.34.08.19

26: **Bandari, V. V., & Lekkala, H. B. (2024).** AI-based operator behavior monitoring and cost optimization using digital traceability in manufacturing systems. *Power System Protection and Control*, 52(1), 38–45. <https://pspac.info/index.php/dlbh/article/view/342>

27: **Venumadhav Vavilala, Shankar Balla (2024, May).** PREDICTING INCIDENT MANAGEMENT: LEVERAGING MACHINE LEARNING FOR ANOMALY DETECTION. *Power System Protection and Control Scopus Q1 Journal. PSPC.* <https://pspac.info/index.php/dlbh/article/view/270>

28: **Manikantha Varaprasad Inakollu, (2024, May),** FINOPS-Driven Cloud optimization models for enterprise applications, Vol. 52 No. 2 (2024): April-June 2024, 99-110, *Power System Protection and Control*, ISSN-1674-3415, URL: <https://pspac.info/index.php/dlbh/article/view/283>, DOI: <https://doi.org/10.46121/pspc.52.2.10>

29: **Naresh Lokiny. (2022).** Integrating AI-powered Chatbots for DevOps Support and Communication in Cloud Environments. *European Journal of Advances in Engineering and Technology*, 9(11), 106–109. <https://doi.org/10.5281/zenodo.13325989>

30: **Naresh Lokiny, (2021),** "Disaster Recovery and Business Continuity Planning in DevOps Cloud with AI", *International Journal of Science and Research (IJSR)*, 10(3), 2024-2027. <https://dx.doi.org/10.21275/SR24724151733>, <https://www.ijsr.net/getabstract.php?paperid=SR24724151733>

31: **Naresh Lokiny, & Ranganath Nandanampati. (2020).** DevSecOps: Integrating Security into DevOps with AI in Cloud. *Journal of Scientific and Engineering Research*, 7(10), 239–242. <https://doi.org/10.5281/zenodo.13348695>

32: **Naresh Lokiny, & Pradip Reddy. (2021).** Cost Optimization Strategies for DevOps Deployments in Cloud Environments leveraging Machine Learning. *European Journal of Advances in Engineering and Technology*, 8(3), 69–72. <https://doi.org/10.5281/zenodo.13325845>

33: **Jayanth Para, 6G Internet Technology** Cyber Threat Notification & Alert System, Vol. 52 No. 4 (2024): October-December 2024, *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/170> , DOI: <https://doi.org/10.46121/pspc.52.4.9>

34: **Jayanth Para, AI Based Cloud Computation** Observational Method & Process, Vol. 51 No. 4 (2023): October-December 2023, *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/171> , DOI: <https://doi.org/10.46121/pspc.51.4.3>

35: **Kaleshwar Aryasomayajula, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES,** Vol. 50 No. 2 (2022): April-June 2022, *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/294>

36: **Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS,** Vol. 50 No. 4 (2022): October-December 2022, *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/364>

37: **Kaleshwar Aryasomayajula, Micro services: “A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments.,** Vol. 51 No. 2 (2023): April-June

2023, Power System Protection and Control, ISSN-1674-3415,
<https://pspac.info/index.php/dlbh/article/view/374>

38: Vikas Suresh Bagora, KERNEL-LEVEL PERFORMANCE OPTIMIZATION FOR CARRIER-GRADE BROADBAND AND HIGH-SPEED NETWORKING SYSTEMS, Vol. 47 No. 1 (2019), Power System Protection and Control, ISSN-1674-3415,
<https://pspac.info/index.php/dlbh/article/view/359>

39: Vikas Suresh Bagora, SCALABLE 128G FIBRE CHANNEL AND ETHERNET CONVERGENCE FOR NEXT-GENERATION DATA CENTER NETWORKS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415,
<https://pspac.info/index.php/dlbh/article/view/362>

40: Vikas Suresh Bagora, SECURE EMBEDDED NETWORKING ARCHITECTURES USING TRUSTED EXECUTION ENVIRONMENTS IN BROADBAND GATEWAYS, Vol. 52 No. 2 (2024): April-June 2024, Power System Protection and Control, ISSN-1674-3415,
<https://pspac.info/index.php/dlbh/article/view/363>

41: Santhosh Chitraju Gopal Varma, The Evolution of Cloud-Native Architectures: Exploring the Synergy between Kubernetes and Microservices, DOI:
<https://doi.org/10.63282/3050-9246.IJETCSIT-V1I4P104>, Vol. 1 No. 4 (2020),
<https://ijetsit.org/index.php/ijetsit/article/view/108>

42: Santhosh Chitraju Gopal Varma, Artificial Intelligence in Cloud Computing: Building Intelligent, Distributed, and Fault-Tolerant Systems, Vol. 3 No. 1 (2022), DOI:
<https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I1P105>,
<https://ijaibdcms.org/index.php/ijaibdcms/article/view/78>

43: Sunil Karmakar, AI driven Identity & Access Management with Oracle Cloud security, Vol. 30 No. 2 (2022): JOCAAA,
<https://eudoxuspress.com/index.php/pub/article/view/4118>