

Post-Quantum Cryptography Migration for Enterprise APIs: A Simulated Hybrid TLS Performance and Risk Evaluation

Neha Verma¹

¹Research Scholar, Department of Computer Science and Engineering, KBNCE, India.

ABSTRACT

This empirical-style article presents a clearly labeled simulated study in Computer Science and Engineering. A simulated enterprise API gateway benchmark compared classical TLS, hybrid classical-plus-ML-KEM key establishment, and post-quantum-heavy configurations under small, medium, and large request profiles. Handshake time, transmitted bytes, CPU utilization, session reuse, memory use, and failure under constrained clients were modeled. The study emphasizes migration engineering and does not claim access to production enterprise systems. Hybrid classical plus ML-KEM increased simulated full-handshake latency by 18% relative to classical TLS but reduced the migration risk associated with an abrupt cryptographic cutover. Session reuse limited the application-level overhead to under 5% for sustained API traffic. Constrained clients experienced the largest memory and payload penalty. The numerical observations, assignments, and measured outcomes are synthetic and intended for methodological illustration only. No real industrial facility, enterprise network, laboratory, construction project, or manufactured component was used. The study demonstrates a realistic research framework aligned with current 2026 research directions while clearly separating simulated evidence from verified literature.

Keywords: Post-Quantum Cryptography, ML-KEM, ML-DSA, Hybrid TLS, Quantum-Safe Security, API Infrastructure

I. Migration Research Questions

The topic addresses a current research direction within Computer Science and Engineering. The central challenge is not only to improve technical performance but also to preserve reliability, interpretability, scalability, and safe operation under real-world constraints.

Because no empirical dataset was supplied, this paper uses a simulated design. The purpose is to create a coherent journal-style analysis with internally consistent methods and results without implying that experiments, fieldwork, or institutional deployment actually occurred.

II. Cryptographic Baselines

A simulated enterprise API gateway benchmark compared classical TLS, hybrid classical-plus-ML-KEM key establishment, and post-quantum-heavy configurations under small, medium, and large request profiles. Handshake time, transmitted bytes, CPU utilization, session reuse, memory use, and failure under constrained clients were modeled. The study emphasizes migration engineering and does not claim access to production enterprise systems.

The modeled variables were selected to reflect the principal engineering mechanisms of the problem. Input ranges were bounded to plausible operating conditions, and controlled noise was introduced so that the results did not become unrealistically deterministic.

III. Hybrid TLS Benchmark Setup

The analysis used separated calibration and evaluation subsets where appropriate. Variables were checked for impossible values, unit inconsistency, duplicated identifiers, and leakage. Sensitivity testing changed one dominant assumption at a time while retaining the remaining conditions.

Robust evaluation requires more than a headline accuracy or efficiency value. The study therefore reports secondary outcomes and sensitivity behavior so that improvements can be interpreted together with their operational cost, failure modes, and uncertainty.

IV. Latency and Payload Results

Hybrid classical plus ML-KEM increased simulated full-handshake latency by 18% relative to classical TLS but reduced the migration risk associated with an abrupt cryptographic cutover. Session reuse limited the application-level overhead to under 5% for sustained API traffic. Constrained clients experienced the largest memory and payload penalty.

Table 1. Simulated primary comparative outcomes.

Configuration	Handshake latency (ms)	Payload (KB)	CPU index	Memory index
Classical TLS	24	7.2	100	100
Hybrid ML-KEM	28.3	10.9	118	124
PQ-heavy	33.6	15.7	134	147




Figure 1. Topic-specific simulated primary result visualization.

V. Migration-Risk Matrix

Table 2. Simulated robustness, subgroup, or sensitivity outcomes.

Client type	Migration difficulty	Quantum-safe benefit
Server API	Low	High
Mobile client	Moderate	High
IoT gateway	High	Moderate
Legacy appliance	Very high	Low




Figure 2. Topic-specific secondary analysis visualization.

Implementation readiness depends on data quality, calibration, maintenance, cybersecurity, human oversight, and compatibility with existing infrastructure. These issues can dominate real-world performance even when the core technology performs strongly in simulation.

VI. Enterprise Transition Strategy

Current literature supports the relevance of the selected technology but also shows persistent gaps between prototype performance and validated deployment. These gaps include generalization, packaging or process integration, long-term reliability, interoperability, standardization, and cost.

A future empirical investigation should define baseline conditions before deployment, reserve independent validation data, document preprocessing and parameter choices, and retain logs of failures or overrides. Where physical safety is involved, fallback behavior should be specified explicitly.

Lifecycle and scalability should also be measured. Hardware aging, software updates, material variation, retraining, environmental exposure, and changing workloads can shift the performance envelope after initial commissioning.

A phased implementation pathway is therefore preferable: controlled validation, limited pilot operation, stress testing, independent replication, and only then wider deployment. Each stage should have predefined acceptance criteria rather than relying on subjective interpretation after results are known.

Economic and lifecycle evaluation should accompany technical benchmarking. Hardware replacement, software updates, retraining, material variability, calibration, safety checks, and operator training can materially change the practical value of a promising method after the initial prototype stage.

Reproducibility would be strengthened by publishing variable definitions, preprocessing rules, simulation assumptions, parameter ranges, evaluation code, and non-sensitive synthetic examples. Where a real dataset cannot be released, sufficient metadata and validation detail should still be provided for independent replication.

VI.A. Cryptographic Inventory

Migration begins with identifying where classical public-key cryptography is embedded. Certificates, TLS termination, VPNs, code signing, device provisioning, backup systems, hardware security modules, and third-party dependencies may all use RSA or elliptic-curve algorithms. An incomplete inventory can leave hidden long-lived vulnerabilities even after visible services are upgraded.

VI.B. Hybrid Migration Strategy

Hybrid key establishment can reduce transition risk by combining classical and post-quantum mechanisms while interoperability matures. The trade-off is additional handshake data, implementation complexity, and the need to verify that both components are composed correctly. Organizations should define when hybrid operation ends and how legacy clients are handled.

VI.C. Certificate and Signature Overhead

Post-quantum signatures can increase certificate and message sizes. This affects handshake fragmentation, memory use, constrained devices, logging systems, and network middleboxes. Benchmarking should therefore include realistic certificate chains and not measure only primitive-level signing or verification time.

VI.D. Crypto-Agility

A successful transition should improve crypto-agility rather than hard-code one algorithm for decades. Algorithm identifiers, key formats, certificate profiles, and protocol negotiation should be modular so that future standards or implementation weaknesses can be addressed without redesigning entire applications.

VI.E. Store-Now-Decrypt-Later Risk

Some data require confidentiality for many years. Attackers may collect encrypted traffic today with the expectation of decrypting it later if sufficiently capable quantum computers become available. Migration priority should therefore account for data lifetime, not merely the expected date of cryptographically relevant quantum computing.

VI.F. Dependency Mapping

Modern applications depend on cryptographic libraries, cloud services, operating systems, load balancers, certificate authorities, and hardware devices. Migration plans should map these dependencies so that one unsupported component does not block the entire quantum-safe transition.

VI.G. Performance Monitoring

Post-migration monitoring should track handshake failures, CPU load, memory pressure, retransmissions, certificate-size effects, and client compatibility. These operational metrics help distinguish cryptographic overhead from unrelated application latency.

VI.H. Transition Roadmap

A staged roadmap can begin with inventory and crypto-agility, continue with hybrid testing on noncritical services, expand to high-value long-lived data, and finish with retirement of vulnerable algorithms once interoperability and standards support are mature.

VI.I. Comparative Evaluation

A mature evaluation should compare the proposed approach with at least one strong contemporary baseline using the same inputs, operating window, and acceptance criteria. This avoids overstating improvement against an artificially weak comparator and makes the engineering significance of the gain easier to judge.

VI.J. Uncertainty Reporting

Engineering decisions should include uncertainty ranges rather than single deterministic values. Bootstrap intervals, scenario bounds, parameter sweeps, and worst-case testing can show whether a reported advantage is stable enough to justify deployment.

VI.K. Lifecycle Monitoring

Performance can change after deployment because hardware ages, software is updated, operating conditions shift, and maintenance alters the system. Long-term monitoring should therefore be built into the design so that drift and degradation are detected before they become safety or reliability problems.

VI.L. Reproducibility and Data Governance

Future empirical work should document data provenance, preprocessing, parameter settings, version control, and evaluation code. Where raw data cannot be released, a detailed data dictionary, summary distributions, and representative synthetic samples can still support independent verification.

VI.M. Protocol-Level Compatibility Testing

Enterprise migration should be tested across real protocol stacks rather than only cryptographic primitives. TLS termination, mutual authentication, API gateways, reverse proxies, certificate validation, session resumption, observability tooling, and network appliances can all react differently to larger keys or signatures. Compatibility testing should therefore include packet fragmentation, timeout behavior, logging systems, certificate automation, and fallback prevention. A successful quantum-safe migration is one that preserves security properties while avoiding silent downgrade to vulnerable classical configurations.

VII. Assumptions and Limitations

The main limitation is that all empirical values are simulated. The reported effect sizes cannot establish real-world effectiveness, causal impact, safety, cost, or generalizability. The model necessarily omits rare failures and site-specific factors.

A second limitation is scenario coverage. Alternative workloads, materials, device parameters, environmental conditions, user behavior, or regulatory constraints could change the ranking of methods. Real validation is required before practical adoption.

VIII. Conclusions

The study demonstrates a distinct empirical framework for a current engineering research trend. The simulated results suggest meaningful potential benefits while also identifying trade-offs and failure modes that would need to be managed in practice. Future work should replace synthetic inputs with measured data, extend validation across environments, and quantify long-term reliability and lifecycle cost.

Declarations

Data availability: All numerical data and empirical results are simulated/illustrative. No real operational dataset was used.

Ethics statement: Not applicable. No human or animal participants were involved.

Funding: No external funding is claimed for this illustrative manuscript.

Conflict of interest: No conflict of interest is claimed for this illustrative manuscript.

References

- [1] Sakwa, C. O., & Li, F. (2026). Survey on post-quantum cryptography implementations and deployment challenges. *Computer Science Review*, 61, 100975. <https://doi.org/10.1016/j.cosrev.2026.100975>
- [2] Thomas, M., Mathews, M. M., Panchami, V., Sebastian, A., & Eswaran, S. (2026). Securing the future: A comprehensive review of post-quantum digital signatures. *Computer Science Review*, 61, 100935. <https://doi.org/10.1016/j.cosrev.2026.100935>
- [3] Liu, Y.-K., & Moody, D. (2024). Post-quantum cryptography and the quantum future of cybersecurity. *Physical Review Applied*, 21, 040501. <https://doi.org/10.1103/PhysRevApplied.21.040501>