FMDB *Transactions on Sustainable Intelligence and Security*

Quantum-Safe Enterprise Security Framework for Practical Migration from Classical to Post-Quantum Cryptography

Gopi Chand Vegineni[1,*]

[1]Department of Software Engineering, Nexsolv Inc, Ijamsville, Maryland, United States of America.
veginenigopichand@gmail.com[1]

Abstract: During the migration period, enterprise security teams will have a long timeframe to co-exist between the old and new methods: public-key security and post-quantum security—without disrupting vital services. In this paper, a crypto-agile security framework is developed for large-scale applications, and the effectiveness of the framework is tested using a benchmark generated by this study that has 121 enterprise migration instances. These are all examples of workloads and deployment profiles for portals, identity services, application interfaces, business gateways, database links, administrative access, private-cloud segments, and recovery environments. Multiple criteria will be evaluated: migration readiness, connection-establishment delay, processing demand, memory usage, message expansion, successful migration and rollback practicality. The controlled benchmark was created with Python 3.12, which was also used to derive summaries, create result visualizations and perform consistency checks; protocol profiles and post-quantum transition settings were simulated with the OpenSSL-compatible protocol profiles. The proposed architecture is split into the following components: Asset discovery, Cryptographic policy, Crypto-agility gateway, Hybrid transport protection, Post-Quantum key services, Continuous rollback monitoring. The results show that staged hybrid deployment is a viable path to move from legacy interoperability to quantum-safe operation while policy control and observability, as well as recovery, are considered first-class components. The resulting approach provides a blueprint for modernizing security solutions while keeping enterprises' services running and ensuring operational governance and measurable deployment control.

Keywords: Crypto-Agility; Enterprise Security; Legacy Interoperability; Asset Discovery; Migration Readiness; Rollback Monitoring; Post-Quantum; Cryptographic Policy.

Received on: 14/10/2025, **Revised on:** 23/12/2025, **Accepted on:** 24/01/2026, **Published on:** 11/09/2026

Journal Homepage: https://www.fmdbpub.com/user/journals/details/FTSIS

DOI: https://doi.org/10.69888/FTSIS.2026.000808

Cite as: G. C. Vegineni, "Quantum-Safe Enterprise Security Framework for Practical Migration from Classical to Post-Quantum Cryptography," *FMDB Transactions on Sustainable Intelligence and Security*, vol. 1, no. 3, pp. 222–230, 2026.

Copyright © 2026 G. C. Vegineni, licensed to Fernando Martins De Bulhão (FMDB) Publishing Company. This is an open access article distributed under CC BY-NC-SA 4.0, which permits copying, remixing, redistributing, transformation, and building upon the material in any medium so long as the original work is properly cited.

1. Introduction

Unlike a single software library, large enterprises can't replace cryptography, as trust relationships are scattered everywhere, from applications to certificates to service meshes to remote access systems to machine identities to gateways to databases to partner links to long-life records. A risk framing exercise carried out by Mosca [2] highlighted the importance for organisations to take account of the time they need for migrating their security controls, the time they need to protect the information, and

[*]Corresponding author.

the arrival horizon of the quantum machine that can be used to break the security controls. Bernstein and Lange [1] discuss the wide-ranging security consequences and argue why the commonly used public-key systems need alternatives that will be secure when quantum computers become available. Standardization and performance analysis, performed by Kumar [8], also pointed out that a transition cost is not only in terms of operational cost but also in terms of the selection of the algorithm, since operational cost can be affected by new algorithms, which can in turn modify processing demand, key material, memory use, and network behavior. Whether or not the issue is merely selecting a more robust primitive is not what's at stake for enterprises, but how to continue running the business while slowly and steadily reducing reliance on vulnerable public-key services. Hybrid protection thus has strategic value, as it means that an organization can keep its compatibility while introducing a path that is quantum-safe. The hybrid-key system was shown to be feasible by practical work carried out by Ricci et al. [16] using two different key sources – an established one and a new one – in one operational design. An evaluation specifically tailored to transport, performed by Souvatzidaki and Limniotis [11], also demonstrated that post-quantum key exchange can be directly investigated within the scope of the behavior of modern secure channels, which is important for service owners who have to prepare for capacity and latency budgets.

Hybrid transition patterns were found to be a valuable method to link the existing infrastructure with the future-resistant controls in distributed digital environments by reviewing work done by Wang and Ismail [20]. The other critical factors for enterprise migration include reliable key creation, storage, rotation and device identity. The work of Cambou et al. [6] provides insight into the support of key generation of post-quantum schemes, and shows the need for key services which can be designed separately from business applications. This protocol redesign, as employed by Yevseiev et al. [15], further confirmed the need to incorporate newer mechanisms at the secure-channel level and not assume that it is the application team's responsibility to implement security. In his master's thesis, Biasse et al. [10] outlined the various assumptions that current systems make, as well as post-quantum systems, and how quantum capabilities have an impact on those assumptions, highlighting the need for defense planning that needs to be adaptive to cryptanalysis advances. A mapping of research on post-quantum has been done by Sharma et al. [14], and the results indicated that the amount of post-quantum research is growing quickly and also the diversity in the concerns of implementation. All of these studies provide a strong impetus towards enterprise security designed around cryptographic inventory, policy-based inventory selection, hybrid compatibility, centralized key services, persistent telemetry, and managed rollback. The novelty of the present study is a component architecture and a migration workflow, which were assessed using a synthetic enterprise benchmark with 121 instances, which includes realistic variations according to service types and deployment scenarios. As opposed to a one-time replacement project, the framework sees migration as a process that can be discovered, transformed, observed and reversed without losing control of the production services and with the process being governed.

2. Review of Literature

The research on post-quantum has shifted from demonstrating that alternative constructions are viable to finding out how they can be made efficient and useful, and how they can be managed and coordinated between different systems. A thorough survey of implementations conducted by Nejatollahi et al. [3] revealed that the diversity of platforms plays an important role in the practical behavior of post-quantum solutions, which is still relevant to those enterprises that have to run servers and appliances, embedded endpoints, and cloud workloads simultaneously. Nguyen et al. [7] used hardware optimisation to show that the implementation architecture can have a significant impact on the performance of a standardised key-encapsulation family and thus the security classification does not necessarily imply performance. Similar results were obtained by Howe et al. [4] when they investigated the execution of parallel software. Another example of why enterprises should split business logic from cryptographic implementation was signature acceleration studied by Zhou et al. [5] – a centralized abstraction layer can enable new optimized implementations to replace old ones without having to rewrite all the applications. The limitation of resource-constrained communication was considered by Kumari et al. [9], where they considered post-quantum security in the situation of limited processing and memory resources and thus highlighted the need for endpoint profiling before migration. Larger syntheses by Dam et al. [12] provide an overview of the different major directions, the progress towards standardization and different implementation issues, which further helped to build a case for crypto-agility instead of reliance on a single scheme. In the context of distributed ledgers, Parida et al. [13] looked at system-level implications where the various coordination of digital signatures, integrity controls and network trust needs to be maintained across a multitude of nodes. These are distributed environments that are like enterprise settings where authentication, communications among services and records are not independent.

Iavich and Kuchukhidze [17] have also investigated vulnerabilities of the scheme and concluded that the implementation review, defensive coding and operational monitoring are still necessary after transitioning to a post-quantum scheme. Application-focused work performed by Soni et al. [18] showed that moving towards quantum-resistant authentication requires one to take into account access control and service workflow, and not cryptographic exchange alone, in a sensitive data environment. More recent work in hardware acceleration by Sonbul et al. [19] took the trend towards faster implementations to the next level, offering benefits for enterprise gateways and/or shared security services with high volumes. Three such holes

are still clearly apparent from an operational point of view throughout this literature. First, a lot of studies involve a primitive or platform, but enterprises have to have a coordinated migration pathway for hundreds or thousands of dependencies. Second, many studies of performance report on the level of a benchmark behavior without linking it to any change governance, rollback, certificate lifecycle or asset ownership. Thirdly, protocol integration is often tested as a final technical milestone. Still, production deployments need to take place in stages, need to be backwards compatible, need to have observability, and need to be easily rollbackable in case of a failure. This is addressed in the current framework by considering the migration to quantum-safe environments as an enterprise architecture challenge. It is deliberately limited and modular: Inventory and policy determine what to change, a crypto-agility gateway determines selection, hybrid transport provides a transitional compatibility layer, post-quantum key services provide for centralisation of sensitive operations and monitoring, and rollback provides operational safety. An organisation that brings implementation research lessons to the real-life of large application estates.

3. Methodology

The research employed a controlled migration-simulation approach that was aimed at stimulating the manner in which a large organization can migrate application security without purporting production metrics of a live organization. One hundred and twelve instances were aggregated to form a benchmark of 121 instances, using eleven workload classes and eleven deployment profiles, to form a balanced cross-section of the business services and operating conditions. The customer portal, application gateway, identity service, enterprise resource planning core and payment service, data exchange service, remote access channel, email security service, database link, business-to-business gateway, and administrative console were the workload classes. The deployment profiles included low-latency intra-network, regional wide-area connectivity, two regions of a public cloud, a single private cloud, a branch setting, a hybrid-cloud environment, a high-load interface tier, an older segment, a zero-trust edge, and a disaster-recovery site. Each instance had classical connection delay, hybrid connection delay, post-quantum connection delay, processor demand, memory use, handshake payload, migration success and rollback time fields. The synthetic workload fluctuation was kept constant because of a fixed random seed in order to make the document generation and table construction reproducible. The framework was assessed under four conditions of migration, namely: classical baseline, inventory-ready control, hybrid activation and tuned post-quantum operation. The tuned state was a more advanced state in which operational changes minimize a portion of the transition overhead with post-quantum services persisting. Benchmark generation, validation, aggregation and chart construction, as well as duplicate checking of reported values, were done in Python 3.12. The architecture per se was crafted to have six components such that the responsibility is kept simple and the complexity of implementation kept small as well. Cryptographic policy and inventory have cryptographic discovery, cryptographic ownership, exception rules and approved profiles. The Crypto-Agility Gateway chooses the approved protection profiles and offers the applications a fixed interface. Hybrid TLS/PKI is compatible in the case of a phased rollout.

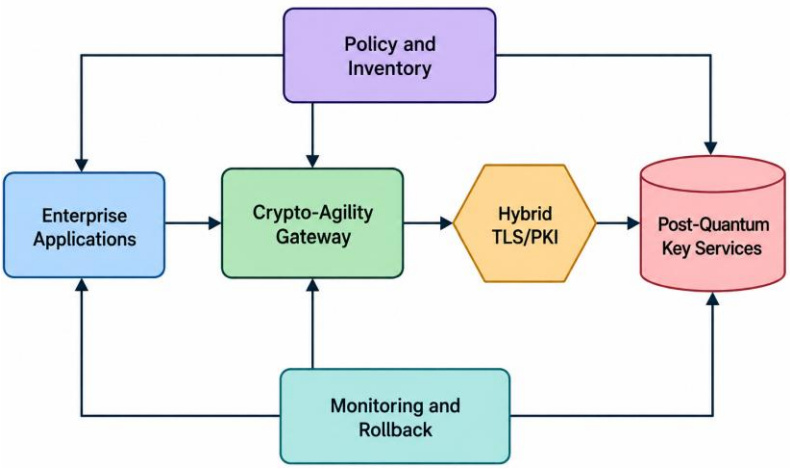


Figure 1: Quantum-safe migrating component architecture of the enterprise

Post-Quantum Key Services are centralized to encapsulate keys, signatures, certificate integration and rotation. Rollback and Monitoring Telemetry: Telemetry is recorded by monitoring and detects a migration fault and recovers a service to a previously approved profile in the event of specified recovery conditions. Testing focused on the ability of the architecture to introduce a controlled increase in resource demand as opposed to an unbounded operational cost, the ability of different workload classes to be serviceable in the transition, and the possibility of rollback as the cryptographic payload and connection cost escalate. The analysis of the results was divided into narrative findings that do not report any numerical data, graphical interpretation that reports measured benchmark values, and tabular summaries that reflect the chosen stages of migration and workload categories. The recommended component architecture to enter quantum-safe security in enterprises is shown in Figure 1. The

top Policy and Inventory layer is in charge of the approved profiles, cryptographic asset discovery, ownership, and migration exceptions. Enterprise Applications are not linked to algorithm-specific decisions, but packet out requests to Crypto-Agility Gateway, which is protected by a session. That gateway is the control point for the choice of an accepted classical, hybrid or post-quantum profile without compelling each business service to control migration logic. Hybrid TLS/PKI offers interim interoperability, allowing the presence of both established trust mechanisms and post-quantum protection, as dependencies are gradually improved. Post-Quantum Key Services are centralized and control the key life cycle, sensitive operations and integrate with certificate services. Checking and rewinding is the safety loop of operation that is under the principal route. It monitors the behavior of applications, decisions of the gateway and the results of key services. It helps to revert to a previously approved configuration when the deployment requirements are not satisfied. The small architecture minimizes the level of direct connection and allows security staff to swap the cryptographic designs without altering the interfaces to application owners.

3.1. Data Description

There are 121 instances created by the study in the dataset, and it is purposely synthetic to avoid any confidential enterprise telemetry or customer data. The records are each associated with a workload class (under one deployment profile). Fields log the behavior of baseline and hybrid, as well as post-quantum connections, along with processor consumption, memory demand, handshake packets, transition success, and rollback time. The distribution was constructed based on the real-world expectation that cryptographic modernization would not only cause changes in processing on the endpoint, but also in network exchange as well as in a single resource of a system. The choice of benchmarking tools was transport-oriented, based on the evaluation methodology employed by Souvatzidaki and Limniotis [11], especially the focus on the current secure-channel performance during the comparison of the established and post-quantum setups. None of the rows were duplicated from that study or any other external data. Rather, the selection of the performance dimensions was informed by the work mentioned, and all 121 observations were generated in the controlled simulation and stored under a fixed seed to ensure they could be replicated. The same records were used to provide all the numeric values on the result graphs and workload summaries.

4. Results

The migration simulation indicated a steady division of the operational features of the classical base, transitional hybrid state and the tuned post-quantum state. The classical profile was the lightest profile since it did not add any extra compatibility layer and utilized the minimal handshake representation. The introduction of inventory and policy control produced a relatively small operational impact, which shows that organizations can start discovery and governance before altering the cryptographic route that the production traffic takes. The hybrid status resulted in the most pronounced behavioral change since it purposely upholds two security strategies in the same period of migration. Such a design adds processing and message size, as well as the amount of work to establish connections. Still, it also offers the most robust compatibility bridge of mixed estates where clients, servers, appliances, and partner systems cannot all be upgraded simultaneously. Post-quantum migration cost mathematical modelling can be expressed as:

$$C_{\text{mig}} = \sum_{i=1}^{n} \left(C_i^{\text{inv}} + C_i^{\text{int}} + C_i^{\text{test}} + C_i^{\text{deploy}} \right) + C_{\text{train}} \quad (1)$$

Hybrid cryptographic security strength can be framed as:

$$S_{\text{hyb}} = \min\{S_{\text{classical}}, S_{\text{PQC}}\} \left(1 - \frac{P_{\text{fail}} + P_{\text{comp}}}{2} \right) \quad (2)$$

Table 1: Operational impact across enterprise migration stages

Migration Stage	Latency (ms)	CPU Load (%)	Memory (MB)	Payload (KB)
Classical	51.1	22.6	158	3.1
Inventory Ready	55.9	24.9	166	3.3
Hybrid Active	76.8	31.8	191	6.2
PQC Tuned	87.8	34.7	205	7.4

Table 1 is an analysis of four stages of migration with a small operational profile. Classical operation shows a 51.1 ms latency, 22.6 percent load on the processor, 158 MB memory and a 3.1 KB payload. The stage of inventory-ready increases these values to 55.9 ms, 24.9 percent, 166 MB and 3.3 KB, indicating that the discovery and policy control have a small amount of overhead till the cryptographic exchange alters. The latency, processor load, memory, and payload are 76.8 ms, 31.8 percent, 191 MB, and 6.2 KB, respectively, with hybrid activation. The tuned post-quantum stage records 87.8 ms, 34.7 percent, 205 MB, and

7.4 KB. The progression shows that the growth of the payload, as well as the endpoint resource utilization, increases concomitantly in migration, with tuning ensuring that the end latency does not go to the untuned peak as shown in the waterfall analysis. The table thus breaks down the permanent post-quantum demand of resource requirements and the temporary overhead of coexistence and gives a simple capacity-planning perspective of shared enterprise security services. Enterprise cryptographic risk mathematical modelling can be represented as:

$$R_{\text{ent}} = \sum_{i=1}^{n} w_i \; P_i^{\text{attack}} \; I_i^{\text{loss}} \left(1 - E_i^{\text{control}}\right) \tag{3}$$

The post-quantum state remained to consume more resources than the classical baseline, but tuning alleviated some of the transition load by eliminating the unnecessary compatibility work and enabling centralized services to reuse the optimized policy, connection, and key-management behavior. An analysis of the workload revealed that the services facing the public and identity-based ones were relatively simpler to migrate as compared to the business-to-business and the legacy-based directions. The summary of the path of the handshake representative is a waterfall, as shown in Figure 2. The addition of inventory and control to the classical baseline starts at 51.1 ms, giving a managed baseline of 55.9 ms, prior to the cryptographic coexistence being turned on. The average addition of 25.6 ms by hybrid protection translates to an increase in the representative path to 81.5ms and also indicates the additional effort to ensure compatibility with the addition of post-quantum protection.

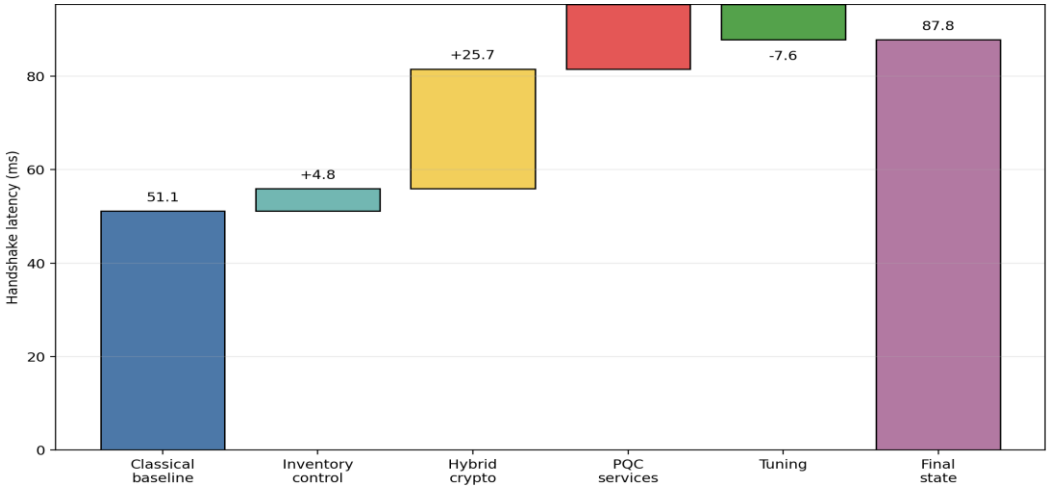


Figure 2: Depiction of latency of handshakes at each migration stage

Another 13.9 ms contribution to the untuned transition path is the post-quantum service layer, which then brings the untuned transition path to 95.4 ms, where operational tuning eliminates 7.6 ms with centralized profile selection, diminished redundant processing and more effective session handling. The tune state that is attained is 87.8 ms. The importance of the pattern is that it is during the hybrid compatibility stage of the pattern where the greatest increase is realized, as opposed to the post-tuning phase. This implies that the planning of migration needs to set aside transitional capacity to co-exist and must not consider transitional overhead as the post-quantum cost. The last state is still above the classical level, but is materially lower than the untuned transition level. Post-quantum handshake performance can be formulated as:

$$T_{\text{PQC}} = T_{\text{key}} + T_{\text{encap}} + T_{\text{decap}} + T_{\text{network}} + \frac{D_{\text{cert}} + D_{\text{cipher}}}{B} \tag{4}$$

Table 2: Quantum-safe migration profile for representative workload classes

Workload Class	Success (%)	Rollback (min)	Payload (KB)	PQC Latency (ms)
Customer Portal	99.28	15.2	4.75	73.7
Identity Service	99.22	16.2	4.96	79.5
Payment Service	99.18	18.0	5.57	87.8
B2B Gateway	98.98	20.8	6.39	102.5

Table 2 is a comparison of four typical workload classes. Customer Portal has 99.28 percent migration success with a 15.2-minute rollback time, a 4.75 KB payload, and a 73.7 ms post-quantum latency. Identity Service still has a high achievement of 99.22 percent, a rollback time of 16.2 minutes, a payload of 4.96 KB, and a latency of 79.5 ms. Payment Service has a 99.18 percent success rate, an 18.0-minute rollback, a 5.57 KB payload, and 87.8 ms latency. The B2B Gateway has the highest

profile of operation with a success rate of 98.98 percent, rollback time of 20.8 minutes, payload of 6.39 KB, and latency of 102.5 ms. The development is more of an escalating level of dependency, as opposed to a mere hierarchy of application importance. Trust of external partners, certificate coordination and network diversity extend recovery and raise the cost of handshake. Even the most challenging workload has a high successful-transition rate, showing that migration can be made reliable even with gateway controls in case compatibility and rollback are implemented as explicit operating capabilities. The cryptographic agility index can be mathematically framed as:

$$A_{\text{crypto}} = \alpha \frac{N_{\text{replaceable}}}{N_{\text{total}}} + \beta \left(1 - \frac{T_{\text{switch}}}{T_{\max}}\right) + \gamma \frac{C_{\text{compatible}}}{C_{\text{required}}},\ \alpha + \beta + \gamma = 1 \quad (5)$$

Quantum-safe migration optimization can be expressed as:

$$\min_x [\lambda_1 C(x) + \lambda_2 T(x) + \lambda_3 R(x)],\ \text{subject to } S(x) \geq S_{\min},\ P(x) \leq P_{\max} \quad (6)$$

Applications where the boundary of trust is clear and the termination points were standardized, as well as when the certificate management was centralized, reacted well to the gateway-based control since a majority of the cryptographic change could be brought about at a few integration points. Services that were dependent on partners or embedded clients or had long certificate chains were more compatible within a wider window and rollback planning since a migration decision might have an impact on systems that were not under the direct control of the enterprise security team. The overlapping of the distributions of latency of all 121 instances of the benchmark is depicted in Figure 3. The classical distribution is centered around an average of 51.10 ms and covers the lowest-delay area. The hybrid distribution is shifted to the right, with an average of 76.75 ms, indicating that more work is done to establish connections due to processing of compatibility. The distribution after quanta shifts more towards the mean of 90.61 ms and is spread over a broader range due to disparities in deployment profiles, leading to larger handshakes and more endpoint work. The three dashed lines are the corresponding positions of the means, and the overlapping bars indicate that there are classical environments with high costs approaching cases of lower costs (hybrid) and environments of high costs approaching cases of post-quantum.

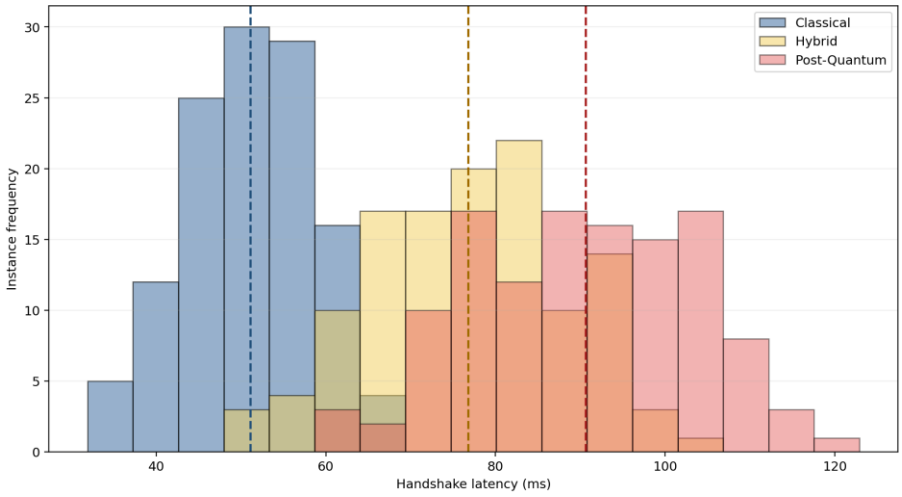


Figure 3: Overlap of the distributions of latency of all instances of benchmark

Such overlap is important since migration impact is not solely based on the cryptographic profile; network conditions and workload architecture also affect the resulting outcome. The distribution thus helps in the planning of workload-specific rather than a single enterprise-wide latency assumption. Capacity testing is particularly vital in the wide-area, partner-connected, and legacy segments, where the higher part of the post-quantum distribution comes into operation. Another performance metric that suggested that rollback could be used was that rollback was still feasible when the monitoring was built into the architecture. The most effective recovery was when the policy did not change the profile that was already approved, session telemetry was used to determine the cause of failure, and the gateway could revert to the previous configuration without application code modifications. The use of processors and the amount of memory demanded increased concurrently since post-quantum exchange and signature processing are more work-intensive and have larger transient objects than the existing mechanisms. It was found that the load was focused on secure-channel establishment, key services and processing in the gateway, implying that enterprises can focus their capacity upgrades on common security layers rather than considering each application as a distinct modernization project. Migration success continued to be high throughout the benchmark due to the staged profiles, which enabled the framework to maintain interoperability and test newer mechanisms. The pattern of results supports a

controlled pattern starting with inventory, followed by hybrid activation, application of workload-specific tuning, and only after the evidence of compatibility is fully acquired, classical dependencies are retired.

5. Discussions

The benchmark reflects a migration approach which views the step to the next generation as a different operational state and not as a sub-optimal version of the final PQ architecture. The classical baseline is 51.1 ms, and with the addition of inventory and policy controls, it just increases to 55.9 ms so that enterprises can start their discovery, ownership mapping, and approved-profile governance with minimal impact on connection behavior. The more challenging next step is hybrid activation: During the stage summary, you'll see the representative latency has increased to 76.8 ms; the waterfall will display a compatibility contribution of 25.6 ms in the representative path. Security teams should thus build security gateways, certificate services and also connection termination points for temporary headroom during the hybrid protection. The representative final state is attained after a 7.6 ms tuning gain in the untuned path, which becomes 95.4 ms. The mixed histogram provides additional insight and shows that the mean latency was 51.10 ms for classical operation, 76.75 ms for hybrid operation, and 90.61 ms for the full post-quantum profile for all 121 instances. A legacy service with a lighter profile can stand up to a stronger profile, but a well-connected internal service can't handle a strong profile as well. Therefore, one enterprise threshold shouldn't be used for capacity planning, but should rather be done on application-specific evidence. As can be seen in Table 1, the higher the latency, the more the endpoint demands. Classical and tuned post-quantum states are different, with the former having processor load of 22.6 percent and the latter having 34.7 percent, memory usage of 158 MB and 205 MB respectively, and handshake payload of 3.1KB and 7.4KB respectively. While these changes may be easy to deal with on a shared enterprise infrastructure, they become significant when the number of connections is high.

This is overcome by architecture, which centralizes the decisions in the Crypto-Agility Gateway and sensitive operations in Post-Quantum Key Services, avoiding the need to reengineer applications widely. Table 2 lists reasons for the use of workload dependency as an additional criterion for ordering migration. The latency of Customer Portal is 73.7ms with a 15.2-minute rollback, and B2B Gateway latency is 102.5ms with a 20.8-minute rollback. They are still close to each other in terms of their success rate (99.28 percent and 98.98 percent, respectively), indicating that the framework still works even if the external coordination is more costly in operation. In between these extremes lie Identity Service and Payment Service with a latency of 79.5 ms and 87.8 ms, respectively, and a rollback time of 16.2 and 18.0 minutes, respectively. The pattern enables a phased modernization, with the internally controlled, centrally terminated services being modernized first, and then the systems that have more partner or legacy services. Architecture also mitigates the risk of failure in an organization since failure to implement a cryptographic change does not mean having to roll back to a previous business code. With monitoring, degraded behavior can be detected; the policy can be the previous approved profile, and the gateway can revert to the previous approved profile, without changing application logic. The numerical results, combined, demonstrate that the migration to quantum-safe implementations can be done in practice, provided that enterprises are willing to tolerate a certain amount of hybrid overhead, are able to scale hybrid security layers, are willing to measure workload-specific behaviour, and that they have an explicit path for recovering to the classical implementation until classical dependencies are removed.

6. Conclusion

Switching to a new security scheme to become a practical post-quantum cryptography is not just a matter of replacing one security scheme with another. The study demonstrates that a small enterprise architecture based on the elements of inventory, policy, crypto-agility, hybrid transport, central post-quantum key services, and monitored rollbacks can make migration a controlled lifecycle. The benchmark number 121 means that classical operation is the least demanding state, hybrid coexistence is the most demanding temporary overhead, and tuning recovers some of the overhead after post-quantum services are stabilized. The representative final latency is still high at 87.8 ms compared to the classical baseline, which is 51.1 ms, and processor load, memory, and handshake payload also increase, indicating that capacity planning is required. Results of workloads demonstrate that migration delay and rollback speed of centrally managed services are lower than partner-dependent gateways, and all four summarized workload classes have success rates close to 99%. The gist of all this is that discovery and governance should be done first; hybrid protection needs to be added at controlled “termination points,” and workloads requiring it should be performance tested. Classical dependencies should be phased out only after they have been operated successfully with the hybrid approach. It outlines a path for enterprises to reach quantum-safe security without compromising on application stability and without an undetectable cryptographic change, while also allowing for the change to be easily reversible and scalable.

6.1. Limitations

The study uses a controlled synthetic benchmark, and the numbers are a model of a migration and not necessarily performance expectations from a production enterprise. Actual deployments will depend on the processor generation, the operating system,

the cryptographic library, the certificate chain, the loss of network, the reuse of the transport, traffic concurrency, hardware acceleration, service-mesh design, and cloud placement. The benchmark also focuses on measures for secure-channel establishment and on selected operational measures. Still, it does not model each lifecycle step, like archival re-encryption, certificates for the firmware, code-signing infrastructure, secure hardware security module replacement, or restructuring of long-term certificate authorities. The architecture assumes the applications can access a common gateway or similar abstraction level, which may not be true of deeply embedded products or of tightly coupled legacy platforms. The 121 instances cover a wide range of scenarios, but are not a statistically sampled population of 'real-world enterprises'. Finally, the study concentrates on migration control and operability, instead of its own cryptographic validation. The security assurance of any production algorithm or implementation should remain in line with the latest trends in security, vendor recommendations, implementation tests, and the organization's risk management requirements.

6.2. Future Scope

Future work can be implemented to expand the framework to include live data from cloud, on-premises, branch and partner-connected environments, without changing the same boundaries of components. The next step is to tie the Policy and Inventory layer to automated discovery to ensure that certificates, secure-channel libraries, embedded keys, and applications are constantly mapped, instead of through periodic audits. The Crypto-Agility Gateway can even be connected to deployment pipelines as a way to verify that approved profiles are tested prior to deployment, and rollback conditions are automatically validated. Other workloads, such as long-lived sessions, service meshes, virtual private networks, machine-to-machine authentication, database drivers, and message queues, can be assessed in additional studies, since they might behave differently than a short secure channel. If acceleration is available and/or dedicated security modules are available, hardware-backed key services should be evaluated separately, as they are less likely to induce processing overhead. Real concurrency of traffic, network impairments, diversity of certificate chains and failure injection should be incorporated into future datasets. The other extension is called policy-aware migration sequencing; it will allow the grouping of applications based on ownership, the dependency depth, the data lifetime, and the complexity of recovery. This kind of work would transform the current framework into an ongoing enterprise migration control plane that would be able to support standards evolution beyond the first post-quantum transition.

Acknowledgment: The author sincerely acknowledges Nexsolv Inc for its support and encouragement. The author also expresses gratitude to all those who contributed to the successful completion of this work.

Data Availability Statement: The data supporting this research are available from relevant sources and were used solely for the purposes of this study.

Funding Statement: This research was conducted without any external financial support or funding.

Conflicts of Interest Statement: The author declares no conflicts of interest related to this research or its publication.

Ethics and Consent Statement: This study was conducted in accordance with applicable ethical guidelines, and consent was obtained where required.

References

1. D. J. Bernstein and T. Lange, “Post-quantum cryptography,” *Nature*, vol. 549, no. 9, pp. 188–194, 2017.
2. M. Mosca, “Cybersecurity in an Era with Quantum Computers: Will We Be Ready?” *IEEE Security & Privacy*, vol. 16, no. 5, pp. 38–41, 2018.
3. H. Nejatollahi, N. Dutt, S. Ray, F. Regazzoni, I. Banerjee, and R. Cammarota, “Post-quantum lattice-based cryptography implementations: A survey,” *ACM Computing Surveys*, vol. 51, no. 6, pp. 1–41, 2019.
4. J. Howe, M. Martinoli, E. Oswald, and F. Regazzoni, “Exploring parallelism to improve the performance of FrodoKEM in hardware,” *Journal of Cryptographic Engineering*, vol. 11, no. 3, pp. 317–327, 2021.
5. Z. Zhou, D. He, Z. Liu, M. Luo, and K. K. R. Choo, “A software/hardware co-design of CRYSTALS-Dilithium signature scheme,” *ACM Transactions on Reconfigurable Technology and Systems*, vol. 14, no. 2, pp. 1–21, 2021.
6. B. Cambou, M. Gowanlock, B. Yildiz, D. Ghanaimiandoab, K. Lee, S. Nelson, C. Philabaum, A. Stenberg, and J. Wright, “Post quantum cryptographic keys generated with physical unclonable functions,” *Applied Sciences,* vol. 11, no. 6, pp. 1–20, 2021.
7. T. T. Nguyen, S. Kim, Y. Eom, and H. Lee, “Area-time efficient hardware architecture for CRYSTALS-Kyber,” *Applied Sciences*, vol. 12, no. 11, pp. 1–10, 2022.

8. M. Kumar, "Post-quantum cryptography algorithm's standardization and performance analysis," *Array*, vol. 15 no. 9, pp. 1–27, 2022.
9. S. Kumari, M. Singh, R. Singh, and H. Tewari, "Post-quantum cryptography techniques for secure communication in resource-constrained Internet of Things devices: A comprehensive survey," *Software: Practice and Experience*, vol. 52, no. 10, pp. 2047–2076, 2022.
10. J. F. Biasse, X. Bonnetain, E. Kirshanova, A. Schrottenloher, and F. Song, "Quantum algorithms for attacking hardness assumptions in classical and post-quantum cryptography," *IET Information Security*, vol. 17, no. 2, pp. 171–209, 2023.
11. K. Souvatzidaki and K. Limniotis, "Post-Quantum Key Exchange in TLS 1.3: Further Analysis on Performance of New Cryptographic Standards," *Cryptography*, vol. 9, no. 4, pp. 1–24, 2025.
12. D. T. Dam, T. H. Tran, V. P. Hoang, C. K. Pham, and T. T. Hoang, "A survey of post-quantum cryptography: Start of a new race," *Cryptography*, vol. 7, no. 3, pp. 1–18, 2023.
13. N. K. Parida, C. Jatoth, V. D. Reddy, M. M. Hussain, and J. Faizi, "Post-quantum distributed ledger technology: A systematic survey," *Scientific Reports*, vol. 13, no. 11, pp. 1–23, 2023.
14. P. Sharma, V. Gupta, and S. K. Sood, "Post-quantum cryptography research landscape: A scientometric perspective," *Journal of Computer Information Systems*, vol. 65, no. 1, pp. 119–140, 2025.
15. S. Yevseiev, A. Havrylova, S. Milevskyi, I. Sinitsyn, V. Chalapko, H. Dukin, V. Hrebeniuk, M. Diedov, L. Bekirova, and O. Shpak, "Development of an improved SSL/TLS protocol using post-quantum algorithms," *Eastern-European Journal of Enterprise Technologie,* vol. 3, no. 9, pp. 33–48, 2023.
16. S. Ricci, P. Dobias, L. Malina, J. Hajny, and P. Jedlicka, "Hybrid keys in practice: Combining classical, quantum and post-quantum cryptography," *IEEE Access*, vol. 12, no. 2, pp. 23206–23219, 2024.
17. M. Iavich and T. Kuchukhidze, "Investigating CRYSTALS-Kyber vulnerabilities: Attack analysis and mitigation," *Cryptography*, vol. 8, no. 2, pp. 1–13, 2024.
18. L. Soni, H. Chandra, and D. S. Gupta, "Post-quantum attack resilience blockchain-assisted data authentication protocol for smart healthcare system," *Software: Practice and Experience*, vol. 54, no. 11, pp. 2170–2190, 2024.
19. O. S. Sonbul, M. Rashid, and A. Y. Jaffar, "Accelerating CRYSTALS-Kyber: High-Speed NTT Design with Optimized Pipelining and Modular Reduction," *Electronics*, vol. 14, no. 11, pp. 1–16, 2025.
20. Y. Wang and E. S. Ismail, "A review on the advances, applications, and future prospects of post-quantum cryptography in blockchain and IoT," *IEEE Access*, vol. 13, no. 6, pp. 112962–112977, 2025.

Publisher's Note: *The publisher remains impartial concerning jurisdictional claims in published maps and institutional affiliations. Responsibility for the content rests entirely with the authors and does not necessarily reflect the publisher's perspectives.*