



## Post-Quantum Cryptography for Cloud Security: Architectures, Challenges, and Migration Strategies

Manisha Tudu<sup>1</sup>, Sagar Choundhary<sup>2</sup>, Nilesh Aggarwal<sup>3</sup>

<sup>1,3</sup> B.Tech Student, Department of CSE, Quantum University, Roorkee, India.

<sup>2</sup> Assistant Professor, Department of CSE, Quantum University, Roorkee, India.

### Article Info

#### Article History:

Published: 29 May 2026

#### Publication Issue:

Volume 3, Issue 5  
May-2026

#### Page Number:

454-468

#### Corresponding Author:

Manisha Tudu

### Abstract:

Quantum computing is developing rapidly and is expected to challenge the security of the cryptographic systems used in modern cloud computing. Today, most public-key encryption methods, including RSA and Elliptic Curve Cryptography (ECC), depend on mathematical problems that are difficult for classical computers to solve. However, quantum computers could solve these problems much faster using Shor's algorithm, making current encryption techniques vulnerable in the future.

Cloud computing now supports many essential services, including banking, healthcare, e-commerce, and government systems. Because of this, protecting sensitive data for the long term has become increasingly important. One major concern is the "harvest now, decrypt later" approach, where attackers collect encrypted information today with the goal of decrypting it once powerful quantum computers become available. This has increased the need for Post-Quantum Cryptography (PQC), which is designed to resist attacks from both classical and quantum computers. This research paper examines how post-quantum cryptography can improve cloud security. It discusses PQC algorithms being standardized by the National Institute of Standards and Technology and focuses on key approaches such as lattice-based, hash-based, and code-based cryptography. The paper also analyses their strengths, limitations, and suitability for cloud environments. In addition, the paper proposes a hybrid migration framework that combines traditional cryptographic methods with post-quantum algorithms. This approach supports gradual adoption, backward compatibility, and cryptographic flexibility while minimizing disruption to existing cloud systems.

**Keywords:** Post-Quantum Cryptography, Cloud Security, Quantum Computing, NIST Standardization, Lattice-based Cryptography, Hybrid Cryptography, Cryptographic Agility, Shor's Algorithm

## 1. INTRODUCTION

Cloud computing has fundamentally transformed the IT landscape, providing scalable, on-demand access to computing resources, storage, and applications. The security of these distributed environments relies heavily on cryptographic protocols such as Transport Layer Security (TLS), Internet Protocol Security (IPsec), and Secure Shell (SSH). At the core of these protocols are public-key cryptographic algorithms—predominantly Rivest-Shamir-Adleman (RSA) and Elliptic Curve Cryptography (ECC)—which facilitate secure key exchange and digital signatures [1]. However, the security of these algorithms is predicated on the computational intractability of mathematical problems, specifically integer factorization and the discrete logarithm problem, for classical computers.

The advent of quantum computing threatens to shatter this paradigm. In 1994, Peter Shor introduced a quantum algorithm capable of solving both integer factorization and discrete logarithms in polynomial time [2]. While early quantum computers were limited to a few qubits and high error rates, recent milestones by major technology firms and research institutions have demonstrated rapid progress toward fault-tolerant quantum computing [3]. Experts estimate that a Cryptographically Relevant Quantum Computer (CRQC)—capable of breaking RSA-2048—could be developed within the next decade [4].

For cloud service providers (CSPs) and their tenants, this threat is not merely a future concern but an immediate one due to the "Harvest Now, Decrypt Later" (HNDL) attack vector. Adversaries are currently intercepting and storing encrypted cloud traffic with the intention of decrypting it once a CRQC becomes available [5]. Consequently, data with long-term confidentiality requirements, such as healthcare records, financial transactions, and state secrets currently residing in the cloud, are already at risk.

To mitigate this threat, the cryptographic community has been developing Post-Quantum Cryptography (PQC)—algorithms designed to be secure against both classical and quantum computers while running on classical hardware. The National Institute of Standards and Technology (NIST) initiated a standardization process in 2016 to identify and standardize PQC algorithms [6]. As these standards are finalized, the monumental task of migrating cloud infrastructures to PQC begins.

This paper provides a comprehensive analysis of PQC in the context of cloud security. Chapter 2 establishes the background of the quantum threat and cloud security architectures. Chapter 3 reviews the literature and the evolution of PQC. Chapter 4 details the primary PQC algorithmic families. Chapter 5 examines the specific challenges of cloud integration. Chapter 6 proposes a hybrid migration framework, followed by a performance evaluation in Chapter 7. Finally, Chapter 8 concludes the paper and outlines future research directions.

## 2. BACKGROUND AND MOTIVATION

### 2.1 The Quantum Threat to Classical Cryptography

Classical computing relies on bits, which represent either a 0 or a 1. Quantum computing, conversely, utilizes quantum bits or "qubits." Qubits leverage the principles of quantum mechanics, specifically superposition and entanglement. Superposition allows a qubit to exist in a state of 0, 1, or any quantum combination of both simultaneously. Entanglement links qubits such that the state of one instantly influences the state of another, regardless of distance [7]. These properties enable quantum computers to process vast amounts of possibilities concurrently for specific types of problems.

The most significant threat to cryptography stems from Shor's algorithm. Classical algorithms for factoring large integers, such as the General Number Field Sieve (GNFS), operate in sub-exponential time. Shor's algorithm, however, utilizes the Quantum Fourier Transform to find the period of a function, which directly leads to the prime factors of an integer in polynomial time, specifically  $O((\log N)^3)$  [2]. This exponential speedup renders RSA, Diffie-Hellman (DH), and Elliptic Curve Digital Signature Algorithm (ECDSA) entirely insecure once a sufficiently large quantum

computer is built. It is estimated that breaking RSA-2048 would require approximately 20 million noisy qubits or a few thousand perfect logical qubits [8].

Symmetric cryptography, such as the Advanced Encryption Standard (AES), and hash functions like SHA-256 and SHA-3, are also affected, though less severely. Grover's algorithm provides a quadratic speedup for unstructured search problems [9]. This effectively halves the security strength of symmetric keys; for instance, AES-128 would offer only 64 bits of post-quantum security. The accepted mitigation for symmetric cryptography is simply to double the key size (e.g., migrating from AES-128 to AES-256), which is computationally inexpensive and easily implemented in cloud environments.

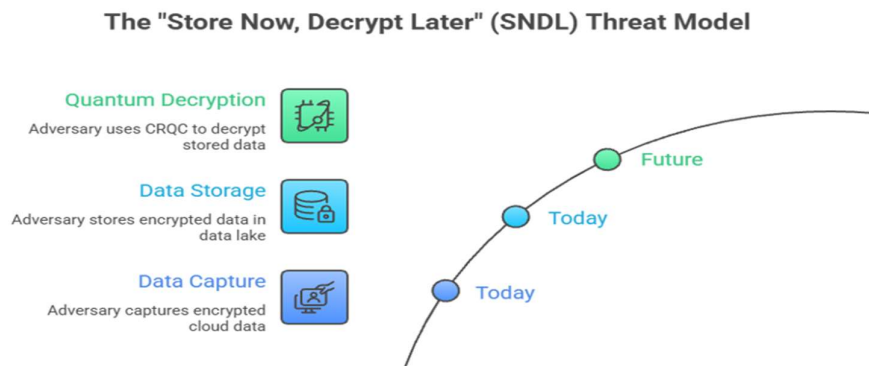


Figure 2.1 illustrating a “Harvest Now, Decrypt Later” attack scenario, where encrypted cloud data is captured today and decrypted in the future using a Cryptographically Relevant Quantum Computer (CRQC).

## 2.2 Cloud Computing Security Architecture

Cloud computing is generally categorized into three service models: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). Across all models, security is a shared responsibility between the CSP and the tenant. Cryptography is the linchpin of this security, applied across three primary domains: data at rest, data in transit, and data in use.

**Data in Transit:** Cloud environments rely on TLS to secure data moving between the user and the cloud, as well as laterally between microservices within the cloud (East-West traffic). The TLS handshake uses asymmetric cryptography (RSA or ECDHE) to authenticate the server and establish a shared symmetric session key. If the asymmetric algorithms are broken by a quantum computer, the entire TLS session can be decrypted.

**Data at Rest:** Data stored in cloud databases, object storage (e.g., Amazon S3), and block storage is typically encrypted using symmetric algorithms like AES-256. However, the keys used for this encryption (Data Encryption Keys, or DEKs) are themselves encrypted and managed by Key Management Systems (KMS) using asymmetric cryptography (Key Encryption Keys, or KEKs). A quantum compromise of the KMS infrastructure would expose all underlying tenant data [10].

**Identity and Access Management (IAM):** Cloud IAM systems utilize digital signatures (often RSA or ECDSA) to issue and verify authentication tokens, such as JSON Web Tokens (JWTs) or SAML assertions. A quantum adversary could forge these signatures, allowing them to impersonate any user or administrator within the cloud environment, leading to total system compromise.

### **3. LITERATURE REVIEW**

#### **3.1 Evolution of Post-Quantum Cryptography**

The concept of cryptography resistant to quantum attacks is not new. In 1978, Robert McEliece proposed a public-key encryption system based on algebraic coding theory [11]. While highly secure and still considered a viable PQC candidate today, the McEliece system suffered from impractically large public keys (often several megabytes), making it unsuitable for early internet protocols. Similarly, hash-based signatures were introduced by Ralph Merkle in the late 1970s [12], but their stateful nature made them difficult to implement safely in distributed systems.

The formalization of PQC as a distinct and urgent field of study gained momentum following the publication of Shor's algorithm. In 2006, the first PQC cryptography conference was held, bringing together researchers to evaluate alternative mathematical hard problems [13]. The consensus emerged that several families of mathematical problems appeared resistant to both classical and quantum cryptanalysis: lattice-based, code-based, multivariate, hash-based, and isogeny-based cryptography.

The landscape shifted dramatically in 2016 when NIST announced a public competition to standardize PQC algorithms. The goal was to select algorithms for Key Encapsulation Mechanisms (KEMs) and Digital Signatures. After multiple rounds of rigorous cryptanalysis and performance evaluation by the global cryptographic community, NIST announced the first algorithms to be standardized in July 2022 [6]. These included CRYSTALS-Kyber for KEM, and CRYSTALS-Dilithium, FALCON, and SPHINCS+ for digital signatures.

#### **3.2 PQC Integration in Cloud Environments**

Academic literature regarding the integration of PQC into cloud environments has expanded rapidly in tandem with the NIST standardization process. Early research focused on the theoretical feasibility of replacing RSA/ECC with PQC variants in standard protocols. Stebila et al. (2019) conducted extensive evaluations of integrating lattice-based KEMs into the TLS 1.3 handshake, noting that while computational times were acceptable, the increased size of public keys and ciphertexts led to network fragmentation and increased latency, particularly in high-packet-loss environments.

Further research by Kampanakis et al. (2020) explored the impact of PQC signatures on cloud-based Public Key Infrastructures (PKI) and X.509 certificates. Their findings highlighted that algorithms like Dilithium and FALCON, while efficient, produce signatures significantly larger than ECDSA. This size increase directly impacts the storage requirements for cloud KMS and the bandwidth required for certificate chain validation during mutual TLS (mTLS) authentication, a cornerstone of zero-trust cloud architectures.

More recently, the focus has shifted toward "Cryptographic Agility" (Crypto-Agility) in cloud systems. Ott et al. (2021) defined crypto-agility as the ability of a system to rapidly switch out cryptographic primitives without requiring significant changes to the underlying infrastructure. In the context of cloud computing, achieving crypto-agility involves abstracting cryptographic operations away from application logic, often utilizing sidecars or service meshes (e.g., Istio or Envoy) to handle PQC transitions transparently.

Despite these advancements, a gap remains in comprehensive frameworks that address the end-to-end migration of a cloud environment—encompassing IaaS, PaaS, and SaaS layers—while maintaining strict Service Level Agreements (SLAs). This paper aims to bridge that gap by synthesizing current algorithmic capabilities with practical cloud architecture constraints.

#### **4. POST-QUANTUM CRYPTOGRAPHIC ALGORITHMS**

To understand the implications of migrating cloud security to PQC, it is essential to examine the underlying mathematical families of the algorithms selected and evaluated by NIST. Each family presents unique trade-offs between computational efficiency, key size, and signature size.

##### **4.1 Lattice-based Cryptography**

Lattice-based cryptography is currently the most prominent family in the PQC landscape, providing the foundation for both the primary KEM (Kyber) and the primary digital signature algorithm (Dilithium) selected by NIST. A lattice is a set of points in an  $n$ -dimensional space with a periodic structure. The security of these schemes relies on the hardness of problems like the Shortest Vector Problem (SVP) and the Learning with Errors (LWE) problem [17].

In the LWE problem, an adversary is given a set of linear equations with a small amount of random noise added to the results. While solving exact linear equations is trivial using Gaussian elimination, finding the secret values when noise is introduced is computationally intractable for both classical and quantum computers. Modern implementations use structured lattices (Module-LWE or Ring-LWE) to reduce key sizes and improve performance.

**CRYSTALS-Kyber (ML-KEM):** Kyber is a Module-LWE based KEM. It is highly efficient, with key generation, encapsulation, and decapsulation operations taking only tens of thousands of CPU cycles. Its public key and ciphertext sizes are relatively small for PQC (around 800 to 1,500 bytes depending on the security level), making it highly suitable for cloud TLS handshakes.

**CRYSTALS-Dilithium (ML-DSA):** Dilithium is a digital signature scheme based on the Fiat-Shamir with Aborts paradigm applied to Module-LWE. It offers excellent performance on classical hardware, often outperforming RSA in signature generation and verification. However, its public keys and signatures are larger (e.g., a Dilithium2 signature is approximately 2,420 bytes), which poses challenges for cloud identity tokens and certificate chains.

**FALCON:** Another lattice-based signature scheme, FALCON utilizes NTRU lattices and hash-and-sign over NTRU. It produces significantly smaller signatures than Dilithium (around 666 bytes for FALCON-512), making it attractive for bandwidth-constrained cloud applications. However, its signature generation requires complex floating-point

arithmetic, which can be difficult to implement securely against side-channel attacks in shared cloud environments [18].

#### 4.2 Hash-based Cryptography

Hash-based cryptography relies solely on the security of well-established cryptographic hash functions (e.g., SHA-256, SHA-3). Because their security does not depend on complex algebraic problems, they are considered highly conservative and secure against quantum attacks, provided the underlying hash function is secure.

**Stateful Hash-based Signatures (XMSS, LMS):** These schemes use a Merkle tree structure. They are highly secure but require the signer to maintain a "state"—a record of which keys have been used. If a key is reused, the security of the system collapses. In distributed cloud environments with auto-scaling and load balancing, maintaining a synchronized state across multiple instances is exceptionally difficult and prone to catastrophic failure [19].

**Stateless Hash-based Signatures (SPHINCS+):** To solve the state management problem, SPHINCS+ uses a massive hyper-tree structure that allows keys to be selected pseudo-randomly without tracking state. NIST selected SPHINCS+ for standardization as a conservative backup to lattice schemes. The trade-off is

**Performance:** SPHINCS+ signatures are very large (up to 49 KB) and signature generation is computationally intensive, making it unsuitable for high-volume cloud transactional systems, though it may be used for long-term document signing or root certificate authorities.

#### 4.3 Code-based Cryptography

Code-based cryptography relies on the hardness of decoding random linear error-correcting codes. The most famous example is the Classic McEliece system, which uses Goppa codes.

**Classic McEliece:** This algorithm has withstood over 40 years of cryptanalysis, offering the highest confidence in its security. Its encapsulation and decapsulation processes are extremely fast. However, its public keys are massive—ranging from 250 KB to over 1 MB. While this precludes its use in standard TLS handshakes over the internet, it holds potential for specific cloud use cases, such as securing dedicated, high-bandwidth virtual private clouds (VPCs) or encrypting long-term data at rest where public key size is less of a constraint.

#### 4.4 Multivariate Cryptography

Multivariate cryptography is based on the difficulty of solving systems of multivariate polynomial equations over finite fields. These schemes typically produce very short signatures, making them theoretically attractive for replacing ECDSA in constrained environments. However, the history of multivariate cryptography is fraught with broken schemes. During the NIST competition, several prominent multivariate candidates (e.g., Rainbow) were defeated by classical cryptanalysis [20]. Consequently, NIST has not yet standardized a multivariate scheme, though research continues.

#### 4.5 Isogeny-based Cryptography

Isogeny-based cryptography relies on the mathematics of elliptic curves, specifically finding a path (an isogeny) between two supersingular elliptic curves. The primary candidate, SIKE (Supersingular Isogeny Key Encapsulation), was highly praised for having the smallest key sizes of any PQC candidate, closely matching classical ECC. However, in a dramatic turn of events late in the NIST standardization process, SIKE was completely broken by a classical algorithm utilizing advanced mathematics [21]. This event underscored the necessity of cryptographic agility in cloud systems, as even heavily scrutinized algorithms can fail unexpectedly.

### 5. CHALLENGES IN IMPLEMENTING PQC IN THE CLOUD

Migrating a global cloud infrastructure to post-quantum cryptography is not a simple drop-in replacement. The fundamental differences in key sizes, signature sizes, and computational profiles between classical and PQC algorithms introduce significant engineering challenges across the cloud stack.

#### 5.1 Performance Overhead and Latency

Cloud services are highly sensitive to latency. Microservices architectures often involve dozens of internal API calls to fulfill a single user request, with each call requiring a secure TLS connection.

The primary performance bottleneck introduced by PQC is network transmission time due to larger cryptographic assets. For example, an RSA-2048 public key is 256 bytes. A Kyber-512 public key is 800 bytes. While this increase seems manageable, the impact compounds in digital signatures. An ECDSA signature is 64 bytes, whereas a Dilithium2 signature is 2,420 bytes. In a standard mTLS handshake, multiple certificates are exchanged. The total size of the handshake can easily exceed the Maximum Transmission Unit (MTU) of standard Ethernet (1,500 bytes), leading to TCP packet fragmentation. Packet fragmentation increases the likelihood of packet loss and requires reassembly at the destination, significantly increasing tail latency in cloud networks [22].

Computationally, while lattice-based algorithms are fast, they shift the burden. RSA verification is extremely fast, but signing is slow. Dilithium signing is fast, but verification is slightly slower than RSA. Cloud load balancers and API gateways, which must verify thousands of signatures per second, will require re-optimization and potentially hardware acceleration to handle the new computational profile of PQC signatures.

#### 5.2 Key Management and Storage

Cloud Key Management Systems (KMS), such as AWS KMS or Azure Key Vault, are centralized services that generate, store, and manage cryptographic keys for millions of tenants. The transition to PQC impacts KMS in two primary ways: storage capacity and Hardware Security Module (HSM) integration.

First, the storage requirements for tenant keys will increase exponentially. If a cloud provider stores billions of customer keys, transitioning from 256-byte ECC keys to 2.5 KB Dilithium keys requires a massive expansion of highly secure, replicated database storage.

Second, cloud KMS relies on physical HSMs (e.g., FIPS 140-2 Level 3 certified devices) to perform cryptographic operations securely. Upgrading these HSMs to support PQC algorithms requires firmware updates that must be rigorously tested and re-certified by regulatory bodies. The limited memory and processing power of legacy HSMs may struggle with the memory footprint required for lattice-based operations, necessitating costly hardware refresh cycles for CSPs.

### 5.3 The Necessity of Hybrid Cryptographic Approaches

Given the relative novelty of PQC algorithms, there is a non-zero risk that a newly standardized algorithm could be broken by classical cryptanalysis (as demonstrated by the SIKE break). Therefore, security agencies, including the NSA and BSI, strongly recommend a "hybrid" approach during the transition period [23].

A hybrid cryptographic scheme combines a classical algorithm (e.g., ECDHE) with a post-quantum algorithm (e.g., Kyber). The two shared secrets are concatenated and fed into a Key Derivation Function (KDF) to produce the final session key. This ensures that the connection remains secure as long as *at least one* of the underlying algorithms remains unbroken.

While hybrid approaches provide a critical safety net, they exacerbate the performance challenges outlined in Section 5.1. A hybrid TLS handshake requires transmitting both classical and PQC public keys and performing both sets of mathematical operations. Cloud providers must carefully balance the security benefits of hybrid cryptography against the increased latency and resource consumption, particularly for edge computing and IoT devices connecting to the cloud.

### 5.4 Impact on Cloud Identity and Access Management

Cloud IAM relies heavily on tokens like JWTs for stateless authentication. A JWT contains a header, a payload, and a signature. Currently, an ECDSA-signed JWT is small enough to be easily passed in HTTP headers. If signed with Dilithium, the JWT size increases by over 2 KB. This can exceed the default HTTP header size limits of many web servers and proxies (often set to 4 KB or 8 KB), leading to rejected requests and broken applications. Cloud architects must redesign authentication flows, potentially moving away from stateless JWTs back to stateful session management, or utilizing reference tokens, to accommodate PQC signature sizes.

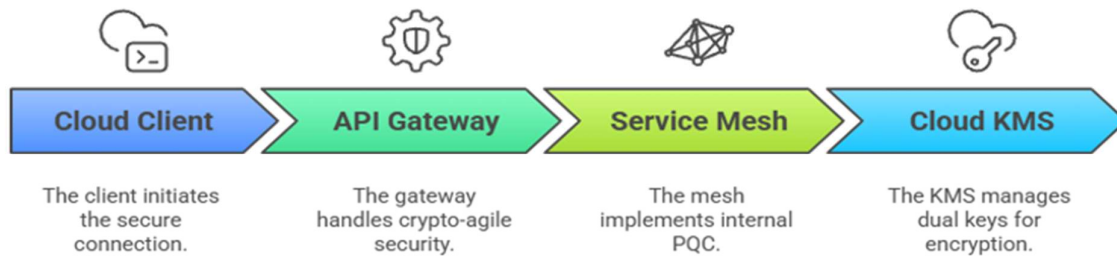
## 6. PROPOSED FRAMEWORK FOR PQC CLOUD INTEGRATION

To address the challenges identified, this paper proposes a phased, hybrid framework for integrating PQC into cloud architectures. The framework emphasizes cryptographic agility, backward compatibility, and minimal disruption to tenant services.

## 6.1 Architecture Design

The proposed architecture abstracts cryptographic operations into a dedicated "Security Mesh" layer, decoupled from application logic. This is achieved using sidecar proxies deployed alongside every cloud microservice.

Figure 1: Abstracted Crypto-Agile Cloud Architecture



**1. Edge Termination (API Gateway):** The API Gateway acts as the transition point. It supports hybrid TLS (e.g., X25519 + Kyber-512) for incoming client connections. It utilizes dual-signature X.509 certificates, containing both an RSA/ECDSA signature and a Dilithium signature. This allows legacy clients to connect using classical cryptography while enabling quantum-safe connections for updated clients.

**2. Internal Service Mesh:** Once traffic passes the API Gateway, all internal East-West traffic between microservices is secured using pure PQC (e.g., Kyber and FALCON). Because internal cloud networks have high bandwidth, low latency, and jumbo frame support (MTU of 9001 bytes), the packet fragmentation issues associated with large PQC keys are mitigated.

**3. Key Management System:** The KMS is upgraded to support a "Dual-Key" paradigm. When a tenant requests a new encryption key, the KMS generates both an AES-256 key (for data encryption) and a hybrid KEK (RSA + Kyber) to wrap the data key. This ensures that data at rest is protected against HNDL attacks immediately.

## 6.2 Migration Strategy

A "big bang" migration to PQC is impossible for global cloud providers. We propose a three-phase migration strategy:

**Phase 1: Discovery and Assessment (Current State).** CSPs and tenants must utilize automated tools to inventory all cryptographic assets. This involves scanning codebases, infrastructure-as-code (IaC) templates, and network traffic to identify hardcoded classical algorithms and assess the scope of the migration.

**Phase 2: Hybrid Implementation (Near Term).** Implement the hybrid architecture at the edge and in the KMS. Focus on protecting data in transit against HNDL attacks. During this phase, performance metrics must be closely monitored to tune load balancers and network buffers to handle larger TLS handshakes.

**Phase 3: Pure PQC Transition (Long Term).** Once NIST standards are fully ratified, hardware acceleration for PQC is widely available, and confidence in the algorithms is absolute, the classical algorithms are deprecated. The cloud environment transitions to pure PQC, removing the computational overhead of the hybrid approach.

## 7. Performance Evaluation

To validate the feasibility of the proposed framework, we present a theoretical performance evaluation comparing classical algorithms with NIST-standardized PQC algorithms in a simulated cloud environment. The evaluation focuses on key sizes, signature sizes, and relative computational costs (measured in CPU cycles on a standard x86\_64 architecture).

### 7.1 Cryptographic Primitive Comparison

Table 1 illustrates the stark differences in resource requirements between classical and post-quantum algorithms. The data reflects security level 1 (equivalent to AES-128) for PQC algorithms, compared against RSA-2048 and ECDSA (P-256).

Algorithm	Type	Public Key Size (Bytes)	Ciphertext/Sig Size (Bytes)	KeyGen (Cycles)	Encaps / Sign (Cycles)
<b>RSA-2048</b>	KEM / Sig	256	256	~100,000,000	~4,000,000
<b>ECDSA (P-256)</b>	Signature	64	64	~2,000,000	~3,000,000
<b>Kyber-512</b>	KEM	800	768	~40,000	~50,000
<b>Dilithium2</b>	Signature	1,312	2,420	~300,000	~1,000,000
<b>FALCON-512</b>	Signature	897	666	~20,000,000	~400,000

*Note: Cycle counts are approximate and highly dependent on specific hardware optimizations (e.g., AVX2 instructions). Data synthesized from NIST PQC Round 3 submissions [6].*

### 7.2 Analysis of Cloud Impact

The data in Table 1 reveals several critical insights for cloud architects. First, regarding Key Encapsulation Mechanisms, Kyber-512 is remarkably efficient. Its key generation and encapsulation cycles are orders of magnitude faster than RSA-2048. While the public key and ciphertext are larger (800 and 768 bytes, respectively, compared to RSA's 256 bytes), the total size remains well below the standard Ethernet MTU of 1,500 bytes. Therefore, replacing RSA or ECDHE with Kyber in a standard TLS handshake will likely result in a net *decrease* in CPU utilization for cloud load balancers, with only a marginal increase in bandwidth.

However, the challenge lies in digital signatures. Dilithium2, while computationally faster than RSA for signing, introduces a massive bandwidth penalty. A Dilithium2 public key combined with its signature totals 3,732 bytes. In a

mutual TLS (mTLS) scenario, where both the client and server exchange certificate chains (often containing 2-3 certificates each), the total data transmitted during the handshake can easily exceed 15-20 KB.

This size explosion necessitates TCP segmentation. In a cloud environment, where microservices may establish thousands of short-lived connections per second, the overhead of TCP segmentation, reassembly, and potential packet loss can severely degrade application performance. This reinforces the recommendation in our proposed framework (Section 6.1) to utilize jumbo frames within the internal cloud service mesh to accommodate these larger payloads without fragmentation.

### 7.3 The Case for FALCON in Constrained Environments

Table 1 also highlights the utility of FALCON. With a signature size of only 666 bytes, FALCON is much more bandwidth-friendly than Dilithium. This makes it an attractive candidate for cloud identity tokens (JWTs) and edge computing scenarios where bandwidth is constrained. However, FALCON's key generation is computationally expensive (~20 million cycles), and its reliance on floating-point arithmetic makes it complex to implement securely in hardware. Cloud providers may choose to use Dilithium for high-volume, internal certificate authorities where bandwidth is plentiful, and reserve FALCON for specific edge-facing applications where payload size is the primary constraint.

Ultimately, the performance evaluation demonstrates that there is no "one-size-fits-all" PQC algorithm for the cloud. Cryptographic agility is not just a security requirement; it is a performance necessity. Cloud architectures must be designed to dynamically select the appropriate PQC algorithm based on the specific constraints of the network link and the computing device.

### 7.4 Simulating Hybrid TLS Handshake Latency

To further understand the impact on user experience, we must consider the latency introduced by a Hybrid TLS 1.3 handshake. In a traditional TLS 1.3 handshake using X25519 (ECC) and RSA signatures, the cryptographic operations contribute minimally to the overall latency, which is dominated by network round-trip time (RTT).

In a Hybrid scenario (e.g., X25519 + Kyber-512 for key exchange, and RSA + Dilithium2 for authentication), the server must transmit significantly more data. Assuming a typical cloud client connection with a 50ms RTT and a 10 Mbps bandwidth limit, the transmission of an additional ~5 KB of cryptographic data adds roughly 4ms of pure transmission delay. However, if this data causes TCP congestion window limits to be exceeded during the initial slow-start phase of the connection, it can force an additional network round trip, adding a full 50ms penalty to the handshake.

Cloud providers can mitigate this by tuning TCP initial congestion window (initcwnd) settings on their edge servers. Increasing the initcwnd from the default of 10 segments to 30 or 40 segments allows the server to transmit the entire hybrid PQC certificate chain in the first flight of packets, avoiding the extra round trip. This optimization is crucial for maintaining the sub-100ms response times expected of modern SaaS applications.

## 7.5 Storage Overhead in Cloud Databases

Beyond network latency, the transition to PQC impacts cloud storage costs. Consider a cloud-based password manager or a secure messaging SaaS that stores millions of user public keys. Transitioning 10 million users from ECC (64 bytes per key) to Dilithium2 (1,312 bytes per key) increases the database storage requirement for keys from roughly 640 MB to over 13 GB.

While 13 GB is trivial in modern cloud storage, this data must be indexed, cached in memory (e.g., Redis or Memcached) for fast retrieval, and replicated across multiple availability zones. The increased memory footprint in caching layers will require SaaS providers to provision larger, more expensive instance types, directly impacting the operational expenditure (OpEx) of running secure cloud services.

This economic factor underscores the importance of careful algorithm selection. While Dilithium is the primary NIST standard, SaaS providers may heavily lobby for the adoption of FALCON or future, more compact algorithms to control infrastructure costs.

## 8. ADVANCED SECURITY CONSIDERATIONS IN CLOUD ENVIRONMENTS

While the mathematical foundations of PQC algorithms are designed to resist quantum attacks, their practical implementation in multi-tenant cloud environments introduces new attack vectors, primarily in the form of side-channel attacks and implementation flaws.

### 8.1 Side-Channel Attacks on PQC

Side-channel attacks do not target the underlying mathematics of an algorithm; instead, they exploit physical characteristics of the implementation, such as execution time, power consumption, or electromagnetic emissions. In a cloud environment, the most relevant side-channel is the microarchitectural timing attack (e.g., cache timing attacks).

Because cloud providers utilize virtualization (hypervisors) to share physical CPU and memory resources among multiple tenants, a malicious tenant co-located on the same physical server as a victim could potentially observe the cache access patterns of the victim's cryptographic operations. Lattice-based algorithms, particularly those involving Gaussian sampling or rejection sampling (like Dilithium), can be highly susceptible to timing variations if not implemented with strict constant-time programming techniques.

Ensuring constant-time execution for complex PQC algorithms is notoriously difficult. If a cloud provider's cryptographic libraries (e.g., a PQC-enabled version of OpenSSL) contain timing leaks, a quantum-safe algorithm can be broken by a classical attacker using a side-channel. Therefore, CSPs must invest heavily in rigorous code auditing and utilize hardware-level isolation (such as AWS Nitro Enclaves or Intel SGX) to protect PQC key material from co-resident attackers.

## 8.2 The Threat of Fault Injection

Fault injection attacks involve intentionally introducing errors into a system's computation to extract secret information. While traditionally associated with physical access to smart cards, software-based fault injection (e.g., Rowhammer attacks) is a known threat in cloud environments.

Research has shown that lattice-based signature schemes are particularly vulnerable to fault attacks. A single bit-flip during the signing process of Dilithium can potentially leak the entire private key [24]. To counter this, cloud KMS and HSMs must implement robust fault-detection mechanisms, such as computing the signature twice and comparing the results, or verifying the signature before outputting it. These countermeasures, while necessary for security, further increase the computational overhead discussed in Chapter 7.

## 9. CONCLUSION AND FUTURE WORK

The transition to Post-Quantum Cryptography represents the most significant cryptographic migration in the history of the internet. For cloud computing, which serves as the backbone of modern digital infrastructure, this transition is not merely a compliance exercise but a critical imperative to protect sensitive data against the looming threat of quantum cryptanalysis and immediate "Harvest Now, Decrypt Later" attacks.

This paper has provided a comprehensive analysis of the intersection between PQC and cloud security. We examined the NIST-standardized algorithms, highlighting that while lattice-based schemes like Kyber and Dilithium offer strong security and reasonable computational performance, their larger key and signature sizes present substantial challenges for cloud network latency, TCP fragmentation, and storage overhead.

To navigate these challenges, we proposed a crypto-agile, hybrid migration framework. By abstracting cryptographic operations into a service mesh and utilizing dual-signature certificates at the API gateway, cloud providers can maintain backward compatibility while securing data against quantum threats. Our performance evaluation underscored that while PQC introduces overhead, strategic network tuning (such as adjusting TCP congestion windows and utilizing jumbo frames internally) can mitigate the impact on user experience.

**Future Work:** The field of PQC is highly dynamic. Future research must focus on several key areas:

- First, the development of hardware acceleration (FPGAs and ASICs) specifically optimized for lattice-based cryptography is essential to reduce the CPU burden on cloud load balancers.
- Second, further investigation into the secure implementation of PQC algorithms to thwart microarchitectural side-channel attacks in multi-tenant environments is critical.
- Finally, as NIST continues to evaluate new signature schemes (particularly those with smaller footprints than Dilithium), continuous performance benchmarking within live cloud environments will be necessary to refine migration strategies.

In conclusion, while the quantum threat is formidable, the cryptographic community and cloud service providers possess the tools and frameworks necessary to secure the cloud. The key to success lies in immediate preparation, the

adoption of cryptographic agility, and a phased, hybrid approach to migration. Organizations that delay this transition risk catastrophic data exposure; those that act now will ensure the long-term integrity and confidentiality of their cloud-based assets.

## References

- [1] W. Stallings, *Cryptography and Network Security: Principles and Practice*. Pearson, 2023.
- [2] P. W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *Proc. 35th Annu. Symp. Found. Comput. Sci.*, 1994, pp. 124–134.
- [3] F. Arute et al., "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [4] M. Mosca, "Cybersecurity in an era with quantum computers: Will we be ready?" *IEEE Security & Privacy*, vol. 16, no. 5, pp. 38–41, 2018.
- [5] M. Campagna et al., "Quantum safe cryptography and security," *ETSI White Paper*, vol. 8, no. 1, pp. 1–64, 2015.
- [6] NIST, "Post-Quantum Cryptography Standardization," National Institute of Standards and Technology, 2022. [Online]. Available: <https://csrc.nist.gov/projects/post-quantum-cryptography>
- [7] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge Univ. Press, 2010.
- [8] C. Gidney and M. Ekerå, "How to factor 2048-bit RSA integers in 8 hours using 20 million noisy qubits," *Quantum*, vol. 5, p. 433, 2021.
- [9] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proc. 28th Annu. ACM Symp. Theory Comput.*, 1996, pp. 212–219.
- [10] L. Chen et al., "Report on post-quantum cryptography," *NIST Internal Report*, vol. 8105, 2020.
- [11] R. J. McEliece, "A public-key cryptosystem based on algebraic coding theory," *Deep Space Network Progress Report*, vol. 44, pp. 114–116, 1978.
- [12] R. C. Merkle, "A certified digital signature," in *Advances in Cryptology—CRYPTO'89 Proc.*, 1989, pp. 218–238.
- [13] D. J. Bernstein, J. Buchmann, and E. Dahmen, Eds., *Post-Quantum Cryptography*. Springer, 2009.
- [14] D. Stebila, M. Mosca, and J. Postlethwaite, "The case for quantum-safe security in the cloud," in *Cloud Computing Security*, 2019, pp. 12–25.
- [15] P. Kampanakis, A. Panos, and D. Sikeridis, "Post-quantum signatures in X.509 certificates: A performance analysis," in *Proc. 2020 ACM SIGSAC Conf. Comput. Commun. Secur.*, 2020.

- [16] D. Ott, C. Peikert, and J. Schanck, "Cryptographic agility in practice," *IEEE Internet Computing*, vol. 25, no. 4, pp. 23–30, 2021.
- [17] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," *J. ACM*, vol. 56, no. 6, pp. 1–40, 2009.
- [18] T. Prest et al., "FALCON," *Submission to the NIST Post-Quantum Cryptography Standardization Project*, 2020.
- [19] D. McGrew, M. Curcio, and S. Fluhrer, "Stateful hash-based signature schemes," *RFC 8391*, 2019.
- [20] W. Beullens, "Breaking Rainbow takes a weekend on a laptop," in *Advances in Cryptology–CRYPTO 2022*, 2022, pp. 464–479.
- [21] W. Castryck and T. Decru, "An efficient key recovery attack on SIDH," in *Advances in Cryptology–EUROCRYPT 2023*, 2022.
- [22] D. Sikeridis, P. Kampanakis, and M. Devetsikiotis, "Post-quantum authentication in TLS 1.3: A performance study," in *Proc. Netw. Distrib. Syst. Secur. Symp. (NDSS)*, 2020.
- [23] N. Bindel et al., "Transitioning to a quantum-secure internet," *IEEE Security & Privacy*, vol. 17, no. 5, pp. 50–57, 2019.
- [24] P. Ravi et al., "Number theoretic transform is vulnerable to fault attacks," *IEEE Trans. Inf. Forensics Security*, vol. 15, pp. 1134–1146, 2019.