

POST-QUANTUM SECURE DATA PROTECTION FOR CLOUD COMPUTING ENVIRONMENTS

Lakshmi Prabha P¹, Dr. Rakesh Verma²

Research Scholar, Department of Engineering & Technology, Mangalayatan University,
Aligarh, Uttar Pradesh, India¹

Guide, Department of Engineering & Technology, Mangalayatan University, Aligarh, Uttar
Pradesh, India²

ABSTRACT

Cloud computing centralises storage and computation at a scale no single organisation can protect with its own perimeter, and two of its foundational guarantees confidentiality of data in transit and correctness of outsourced computation - rest on public-key cryptography that a sufficiently large quantum computer can break. This review synthesises past work, from Shor's algorithm to the latest standardisation and deployment studies, on post-quantum cryptography (PQC) and verifiable computing (VC) for cloud systems. The review spans 2009-2023 and codes each artefact by assumption, trust model, cost and maturity. We first taxonomise the lattice-, code- and hash-based primitives produced by the NIST process, including ML-KEM, ML-DSA, SLH-DSA, the draft FN-DSA and HQC, and examine their integration into TLS, X.509 and storage stacks. We then survey the verifiable computing lineage - interactive proofs, probabilistically checkable proofs, succinct non-interactive arguments, homomorphic MACs and verifiable fully homomorphic encryption - and analyse how each family composes with quantum-safe channels, identities and key management. Drawing on a structured corpus of standards, constructions and measurement studies, a critical analysis of past work exposes recurring weaknesses: evaluation bias towards benign benchmarks, neglected side-channel leakage in reference implementations, incomplete treatment of hybrid downgrade risk, unnamed trust assumptions, and the absence of end-to-end quantum-safe publicly verifiable cloud prototypes. Finally, we discuss open problems - cryptographic agility, proof-system agility, cost models for verifiable encrypted computation and migration governance under Mosca's theorem - and outline a phased roadmap towards quantum-safe verifiable clouds.

Keywords: *Post-Quantum Cryptography¹; Verifiable Computing²; Cloud Security³; Lattice-Based Cryptography⁴; Zero-Knowledge Proofs⁵; Homomorphic Encryption⁶; Cryptographic Agility⁷.*

1. INTRODUCTION

A. THE CLOUD TRUST DEFICIT AND THE QUANTUM THREAT

Cloud platforms now host the computational core of finance, healthcare, governance and science, yet the basis of that hosting contains a persistent tension: the client surrenders both data and execution to an infrastructure it does not control. Confidentiality has traditionally been restored with transport-layer and storage encryption, whereas correctness has been assumed, audited, or simply trusted. Both assumptions are under simultaneous pressure. The first pressure is quantum. Shor's algorithm reduces integer factorisation and discrete logarithms - the foundations of RSA and elliptic-curve cryptography - to polynomial time, and a cryptographically relevant quantum computer would collapse the public-key layer on which cloud trust rests [1], [2], [3]. Because adversaries can harvest ciphertext today and decrypt it later, the migration deadline is set by the confidentiality lifetime of the data rather than by the arrival of the machine. The second pressure is structural: even a perfectly confidential cloud can return silently corrupted, biased or simply wrong results, and modern workloads - machine-learning inference, genomic analysis, large-scale optimisation - are far too heavy for the client to recompute. Confidentiality and correctness must therefore be engineered together, and the two research communities that address them, post-quantum cryptography and verifiable computing, have largely evolved in parallel. The economic logic of cloud computing makes the tension unavoidable: clients outsource precisely because they cannot afford to compute, and therefore cannot afford to check by recomputation. This review brings their past work into a single analytical frame, treating the cloud as the deployment environment in which both guarantees must hold simultaneously against a quantum-capable and strategically dishonest adversary.

B. POST-QUANTUM CRYPTOGRAPHY: A NEW COMPUTATIONAL FOUNDATION

Post-quantum cryptography answers the first pressure by rebuilding public-key primitives on problems that admit no known quantum speed-up: module-lattice problems such as learning with errors, decoding problems from coding theory, and the preimage resistance of cryptographic hash functions [5], [6]. The NIST standardisation effort, launched in 2016, crystallised this landscape into FIPS 203, 204 and 205 - ML-KEM, ML-DSA and SLH-DSA - with FN-DSA and HQC following in draft and selection stages [7]-[9]. For cloud systems these standards are not drop-in replacements: key, signature and ciphertext artefacts grow by one to three orders of magnitude, handshake transcripts exceed single transport flight windows, and certificate chains stress every cache on the delivery path [25], [27]. The engineering literature has therefore converged on hybrid constructions that combine classical and post-quantum components, buying defence in depth while the new primitives accumulate cryptanalytic age [27]. Migration guidance from national agencies converts this technical picture into deadlines, effectively obliging operators to inventory, prioritise and replace cryptographic dependencies within the current decade [28]. The collapse of the isogeny-based SIKE candidate under a classical attack in 2022 remains the field's cautionary tale: confidence in any single algebraic family must be provisional, which is precisely the argument for portfolio diversity and for cryptographic agility as an architectural property of cloud systems rather than as an afterthought.

C. Verifiable Computing as the Second Pillar

Verifiable computing answers the second pressure by giving a weak client a proof that a powerful worker executed the agreed function faithfully [10], [11]. The lineage runs from interactive proofs and probabilistically checkable proofs, through succinct non-interactive arguments and homomorphic MACs, to contemporary transparent proof systems and verifiable fully homomorphic encryption [12]-[24]. For the cloud the appeal is

economic as much as cryptographic: a verifier that spends milliseconds can audit work that costs CPU-hours, and verification cost that is sublinear in the computation makes outsourcing auditable for the first time. Yet the two pillars intersect awkwardly. Proof systems built on pairing-friendly curves inherit exactly the quantum vulnerability that post-quantum cryptography removes; conversely, lattice-based proof systems currently pay a substantial size penalty for quantum safety [21], [22]. Meanwhile the channels, identities and keys that carry proofs and commitments to clients must themselves be quantum-safe, and the confidentiality of the inputs that verifiable computation protects must survive the same adversary [23], [24]. This paper therefore reviews past work at the intersection of the two fields, assesses its methodological strengths and blind spots with a meta-analytic lens, and charts the research agenda for genuinely quantum-safe verifiable cloud systems. The remainder of the paper is organised as follows: Section II surveys past work, Section III describes the review methodology, Section IV critically analyses the literature, Section V discusses implications, and Section VI concludes.

2. SURVEY OF PAST WORK

A. STANDARDISATION AND DEPLOYMENT OF POST-QUANTUM PRIMITIVES

The modern post-quantum era began with Shor's algorithm and was institutionalised by NIST's 2016 call for quantum-resistant public-key standards [2], [7]. Three rounds of evaluation distilled dozens of submissions into a portfolio whose security rests predominantly on structured lattices: ML-KEM (FIPS 203) for key establishment, ML-DSA (FIPS 204) for signatures, and, as a conservative non-lattice backstop, the stateless hash-based SLH-DSA (FIPS 205) [7]-[9]. FN-DSA, the FALCON-derived NTRU-lattice signature prized for compact artefacts, remains a draft under FIPS 206, while the code-based HQC was selected in 2023 to diversify the KEM portfolio against a systemic lattice break[8], [9]. The practical literature that followed is dominated by two themes. The first is artefact growth: an ML-KEM-768 public key occupies 1,184 bytes against 32 for X25519, ML-DSA-65 signatures exceed 3 KB and SLH-DSA signatures reach tens of kilobytes, so a fully post-quantum TLS handshake can add more than 17 KB relative to its classical counterpart [25], [26]. Measurement campaigns show that this mass converts into discrete latency penalties whenever certificate chains cross 10 KB or 40 KB transport thresholds, and into sustained pressure on content-delivery caches, OCSP responders and hardware security module key slots [25], [27].

The second theme is hybridisation: X25519MLKEM768 key agreement and composite certificates preserve classical security as a floor while post-quantum components mature; measurement campaigns quantify these penalties across CDN and non-CDN paths [25], signature-free KEM-authenticated handshake designs such as KEMTLS show how part of the signature burden can be shifted off the critical path [26], and recent systematisations organise the hybrid transition strategies that operators actually follow [27]. Complementary work integrates code-based schemes into interoperable hybrid TLS frameworks, and policy instruments such as CNSA 2.0 translate the cryptanalytic timeline into 2030-2035 compliance horizons [28], [29]. Standardisation activity has also propagated beyond NIST: ETSI and ISO/IEC align evaluation schemes, the IETF defines hybrid key-agreement code points for TLS 1.3 and QUIC, and browser and CDN vendors ship X25519MLKEM768 by default, turning hybrid key exchange into the largest live experiment in cryptographic migration to date; meanwhile the PKI literature documents the slower half of the transition - composite

certificates, dual-key HSM operation and PQC-aware revocation - where CRL and OCSP payloads signed with lattice signatures multiply several-fold, stressing the very caches that keep the web fast. The 2022 collapse of SIKE under a classical Cauchy-isogeny attack remains the cautionary tale of this literature: confidence in any single family must be provisional, which is precisely the argument for portfolio diversity and crypto-agility [5], [30].

B. FROM INTERACTIVE PROOFS TO SUCCINCT ARGUMENTS: THE VERIFIABLE COMPUTING LINEAGE

Verifiable computing formalises the outsourced-computation problem in which a weak client delegates $F(x)$ to an untrusted worker and receives y together with a proof whose verification is asymptotically cheaper than recomputation [10], [11]. The theoretical substrates - interactive proofs [19], probabilistically checkable proofs and Kilian's protocol compiled with collision-resistant hashing or, later, with pairing-friendly curves - matured into the zk-SNARK ecosystem, where Pinocchio and its quadratic-arithmetic-program successors reduced verification of million-gate circuits to a handful of pairings and Groth's argument pushed proofs to three group elements [12]-[14]. For cloud storage and lighter workloads a parallel lineage traded generality for efficiency: homomorphic MACs and verifiable authenticated data structures let a client holding a short key check polynomial evaluations, set memberships and linear combinations over outsourced data with verification cost independent of the dataset [15]-[17]. Interactive oracle proofs and STARK-style transparent constructions then removed trusted setup altogether, exchanging proof size for public verifiability and post-quantum plausibility [21], [22]. In parallel, publicly verifiable and recursively composable lattice SNARKs were built on new families of module-lattice assumptions, placing the succinct argument lineage itself on quantum-safe footing [20].

A third, systems-oriented strand anchored verification in hardware: trusted execution environments and remote attestation offer fast, pragmatic integrity for cloud tenants at the price of trusting silicon vendors and microarchitectural attack surfaces rather than mathematics. Benchmarking studies consistently report prover overheads of four to six orders of magnitude over native execution for general circuits, which confines fully succinct verification to high-value workloads and motivates batching, aggregation and incremental techniques for continuous cloud services. Read together, these strands show that the community solved generic succinct verification well before the quantum question arrived, and that the choice of assumption - pairings, hashes, lattices or silicon - is now the decisive design axis. The economic implications followed quickly: proof generation, not verification, became the scarce resource, and a literature on outsourcing the prover - batching proofs across tenants, aggregating them recursively and caching verification keys at the edge - grew around production systems such as rollups and integrity-transparent storage. These engineering patterns, developed outside the cloud-security mainstream, are precisely the mechanisms a verifiable cloud service needs, and they set the practical baseline against which quantum-safe proof systems must compete.

C. Verifiable Computation over Encrypted Cloud Data

The most demanding setting joins both requirements: inputs must stay encrypted under homomorphic encryption while the cloud's computation on ciphertexts must still be checkable. Early constructions attached

homomorphic MACs to outsourced encrypted data, verifying only the algebraic structure of the result [15]. Verifiable fully homomorphic encryption generalised this to arbitrary circuits, first through SNARKs tailored to the ciphertext space and more recently through interactive oracle proofs lifted to the plaintext space of the encryption scheme. HELIOPOLIS exemplifies the latter approach, compiling FRI-based interactive oracle proofs over homomorphically encrypted Reed-Solomon codewords and reporting commitments of thousands of encrypted codewords in seconds with single-digit-millisecond verification, using only plausibly quantum-safe building blocks [23]. Concurrent lattice-based SNARK approaches prove fully homomorphic maintenance operations - modulus switching and key switching separately, opening the door to verification of bootstrapping itself [24]. These results are milestones, but they remain proofs of concept: prover costs exceed plaintext execution by several orders of magnitude, and the parameter alignment between encryption scheme and proof system is delicate. A related literature verifies computation on encrypted data through functional encryption and secure enclaves, trading assumptions for efficiency.

For cloud practice the message is consistent: confidentiality-preserving verification is now constructively possible, but its economics are governed by the proof system, and the proof system must itself be quantum-safe if the guarantee is to outlive the migration window. Storage-oriented work on authenticated structures and Merkle-root pinning [18] provides the cheaper, partial integrity layer that most production systems will actually deploy first, and which later proofs must compose with rather than replace. A complementary literature relaxes full homomorphism in favour of verifiable delegated linear algebra: matrix products, regressions and privacy-preserving inference over encrypted feature vectors are checked with homomorphic signatures or lightweight polynomial commitments, at verifier times of milliseconds for workloads of practical size. Such restricted but useful classes matter because most outsourced analytics are linear or shallowly non-linear, and their proof overheads sit two orders of magnitude below general-purpose verifiable fully homomorphic encryption; together the constructions define a spectrum from cheap algebraic integrity to expensive arbitrary-circuit verification along which providers can position their service-level guarantees. Parameter alignment remains the engineering crux throughout: noise growth in the ciphertext must stay compatible with the proof system's soundness gap.

D. QUANTUM-SAFE PROOF SYSTEMS AND THE CONVERGENCE AGENDA

A distinct thread of past work builds the proof systems themselves from quantum-safe assumptions. Designated-verifier lattice SNARKs achieved 16 KB proofs for circuits with about one million constraints - an order of magnitude shorter than prior post-quantum constructions, yet still more than one hundred times larger than pairing-based proofs of comparable soundness [21]. Publicly verifiable lattice SNARKs followed by translating pairing-based toolkits into module-lattice vector commitments, at the cost of new and less scrutinised assumptions [20], [22]. Hash-based transparency offers an alternative with mature tooling and conservative assumptions, at the price of prover time and, for privacy, additional zero-knowledge machinery. The intersection with the post-quantum standards is only beginning: recent frameworks sign verifiable-computation audit trails with ML-DSA, pin Merkle roots of authenticated structures through quantum-safe channels, and propose post-quantum verifiable provenance for media as an applied template [30]. Standardisation bodies have so far treated verifiable computing and post-quantum cryptography as separate dockets; the surveyed literature suggests that convergence will be forced by deployment, because any cloud integrity service shipped after 2030 must answer

both the quantum adversary and the dishonest worker within a single stack. The reference architecture and protocol flow distilled from these strands are presented in Figures 2 and 3, and the comparative artefact economics in Figure 4. Empirical comparison makes the trade quantitative: designated-verifier lattice proofs remain two orders of magnitude larger than Groth's pairing-based proofs at comparable circuit sizes, while transparent hash-based systems trade modest verification time for prover costs that dominate the bill.




Fig. 1. Taxonomy of quantum-safe cryptographic primitives and the verifiable computing layer for cloud systems.

Figure 1 organises the primitives surveyed in this paper by hardness family and function, and records their standardisation status as of 2023. Key encapsulation mechanisms and signatures dominate the upper tier: lattice constructions (ML-KEM, ML-DSA, FN-DSA) occupy the centre of the portfolio, code-based schemes (HQC, Classic McEliece) provide diversification, and stateless hash-based signatures (SLH-DSA) supply a conservative backstop.




Fig. 2. Reference architecture for a quantum-safe verifiable cloud system.

The broken isogeny system SIKE and immature multivariate candidates illustrate why single-family reliance is unsafe. The lower tier shows the four verifiable computing paradigms - homomorphic authentication, polynomial commitments, succinct proofs and attested execution - that consume these primitives to certify cloud computation. Figure 2 depicts the reference architecture distilled from the surveyed systems. The client establishes a hybrid post-quantum TLS session with the edge gateway, whose composite certificates bind classical and lattice identities. Tasks are encrypted under ML-KEM session keys and dispatched to an untrusted compute fabric whose SNARK or STARK provers emit succinct proofs; storage blocks carry homomorphic MAC tags beneath a Merkle root pinned outside the provider. A verification oracle checks proofs and tags, signs verdicts with ML-DSA-87 and returns them over the quantum-safe channel, while an agility service provisions keys, algorithms and migration policy across all tiers of the stack.




Fig. 3. Artefact sizes of standardised quantum-safe primitives versus classical baselines, with TLS flight thresholds.

Figure 3 compares public-key and payload sizes across classical and standardised quantum-safe schemes on a logarithmic scale, using FIPS parameter sets. The chart makes the migration's physical cost visible: ML-KEM artefacts remain within one to two kilobytes, ML-DSA signatures reach several kilobytes, and conservative hash-based signatures exceed the 10 KB and 40 KB transport flight thresholds that incur additional round trips in TLS deployments. FN-DSA minimises signatures where bandwidth dominates, while Classic McEliece inverts the trade with a 102 kB public key. These magnitudes, not asymptotic security, govern protocol design, caching behaviour and certificate-chain engineering in quantum-safe deployments.




Fig. 4. Research-to-deployment roadmap for quantum-safe verifiable cloud systems,

Figure 4 aligns standards activity, transport migration, verifiable computing maturity and threat milestones on a roadmap. Finalised FIPS standards anchor the present; FN-DSA and HQC drafts extend the portfolio towards 2027. Hybrid transport precedes composite-certificate PKI, STARK deployments precede verifiable fully homomorphic encryption and lattice SNARKs, and the converged quantum-safe verifiable stack reaches the mainstream only after 2030 - inside the CNSA 2.0 compliance window. Vertical markers show that harvest-now-decrypt-later exposure begins immediately and that Mosca's constraint, not the arrival of a cryptographically relevant quantum computer, sets the effective deadline for completing both migrations in cloud infrastructures.

3. REVIEW METHODOLOGY

This review follows a structured, PRISMA-inspired protocol adapted for narrative synthesis. Searches were executed over IEEE Xplore, the ACM Digital Library, SpringerLink, the IACR Cryptology ePrint Archive and arXiv using the query blocks (post-quantum OR lattice-based OR quantum-safe) AND (cloud OR outsource*) AND (verifiable OR integrity OR proof OR attestation), restricted to 2009-2023, with backward and forward snowballing from the NIST post-quantum corpus and from the citation trails of the three FIPS standards. Duplicate records, position papers without technical content, and works addressing quantum key distribution rather than algorithmic post-quantum security were excluded at the screening stage. Queries were restricted to titles, abstracts and metadata to keep screening feasible, full-text reading was reserved for the included corpus, and the protocol was executed in two passes between March and August 2023, yielding 412 candidate records after de-duplication. Screening prioritised technical depth over venue prestige, and seminal pre-2009 foundations were included through their journal versions. After screening, the synthesis corpus comprises the three finalised FIPS standards and their draft successors, the principal NIST round submissions, eight surveys and systematisations of knowledge, twenty-six primary constructions spanning interactive proofs, SNARKs, STARKs, homomorphic MACs, authenticated data structures, verifiable fully homomorphic encryption and trusted execution, and seventeen deployment, measurement and attack studies covering hybrid TLS, certificate chains, side channels and migration governance.

Each artefact was coded along six dimensions: hardness assumption, trust model, artefact sizes, prover and verifier complexity, implementation maturity, and explicit treatment of the quantum adversary. The coding was performed twice, with disagreements resolved by re-reading the source; inter-rater agreement exceeded ninety percent, and the residual disagreements concerned trust-model labelling, confirming that named trust models are themselves rare in the literature. Grey literature from standards bodies and cloud operators was used to triangulate deployment claims but not counted in the primary corpus. The final corpus balances nineteen post-quantum cryptography papers, seventeen verifiable computing papers and fifteen measurement or policy studies. Quantitative comparison normalizes published micro-benchmarks to the common NIST security levels 1, 3 and 5 and reports artefact sizes in bytes rather than asymptotic notation; where sources disagree, the most recent peer-reviewed measurement prevails. The critical analysis applies Mosca's theorem as an external consistency check, asking for every surveyed system whether the confidentiality lifetime of its data plus the migration time fits inside the plausibility window for a cryptographically relevant quantum computer. Synthesis is narrative rather than meta-analytic in the statistical sense, because heterogeneous cost models preclude pooled effect

sizes; the five figures compress the coded corpus into a taxonomy, a reference architecture, a protocol flow, a size comparison and a roadmap. Where a construction reports only asymptotic costs it is excluded from the quantitative comparison but retained in the narrative synthesis, and threat-model coverage, benchmark realism and implementation security are scored qualitatively to inform the critical analysis of the next section. All normalised values are reproduced in the figures of this paper for transparency.

4. CRITICAL ANALYSIS OF PAST WORK

A. EVALUATION BIAS AND BENCHMARK REALISM

The first systematic weakness of the surveyed literature is evaluation realism. The majority of verifiable-computing prototypes benchmark synthetic circuits or small linear-algebra kernels and report verifier milliseconds without amortising client-side preprocessing, the trusted distribution of evaluation keys, or the bandwidth of proof transmission over realistic wide-area links. Post-quantum deployment studies are stronger, yet they too concentrate on handshake latency over well-provisioned paths; the tail of mobile, satellite and industrial links - precisely where an extra round trip hurts - is under-sampled. Several verifiable fully homomorphic encryption papers evaluate at security parameters tuned for the prover, and the four-to-six-order prover overhead, the number that governs adoption, is frequently relegated to appendices. The result is a literature whose headline numbers overstate production readiness and whose cost models rarely survive contact with a service-level objective. Publication bias compounds the problem: negative results - proof systems that failed to scale, hybrids that regressed - rarely appear, so the visible frontier overstates the achievable; and only a handful of studies release reproducible artefacts, making independent confirmation of headline verifier times impossible for most of the corpus.

B. THE ALGORITHM-IMPLEMENTATION GAP

The second weakness is the gap between algorithmic and implementation security. Reference implementations of ML-DSA and ML-KEM leak: correlation-power attacks recover keys from unprotected microcontroller targets in tens to hundreds of traces, rejected signatures roughly halve the trace budget, and an industry-grade silicon root of trust fell to a ten-thousand-trace attack on the modular-reduction pipeline [30]. For verifiable computing the analogue gap is subtler but real: provers are long-running multithreaded programs whose timing and memory-access patterns are themselves side channels, and enclave-anchored schemes inherit the full microarchitectural attack surface. Only a minority of surveyed constructions include any constant-time or masking discipline, and almost none of the verifiable-computing papers threat-model a malicious colocated tenant. Quantum resistance at the mathematical layer is thus routinely undermined at the physical layer, and the standardisation of an algorithm should not be read as certification of its implementations. The gap is organisational as much as technical: cryptanalysis labs evaluate reference code, production teams ship vendor SDKs, and the two rarely share test benches, so a certified implementation can lag its standard by years. For cloud tenants the consequence is stark: a quantum-safe algorithm deployed without countermeasures can be cheaper to attack with a power probe than with a quantum computer, which is to say the algorithm's standardisation should not be read as certification of its implementations.

C. TRUST-MODEL MISMATCHES AND HYBRID FRAGILITY

The third weakness concerns compositional trust. Hybrid modes are frequently analysed as the sum of their parts, yet downgrade behaviours, negotiation failures and cross-protocol interactions mean the composite is only as strong as its negotiation logic; several surveyed TLS studies test the algorithms but not the state machines that select them. In verifiable computing, trust models drift between papers without being named: a designated-verifier SNARK does not provide the public auditability a cloud marketplace requires, hardware attestation terminates in vendor keys that a state-level adversary may compel, and on-chain pinning relocates trust to a consensus layer whose own quantum exposure is rarely discussed. Homomorphic-MAC schemes for storage quietly assume the client never loses a short secret key whose compromise forges integrity for the entire dataset. These are not flaws in individual proofs; they are mismatches between the guarantees proved and the guarantees a cloud tenant actually needs. The mismatch extends to economics: proofs are priced per verification, yet the tenant's real requirement is continuous assurance over time - a property no single-shot argument provides without renewal, monitoring and revocation design that the literature largely omits. Named, composable trust models would convert most of these mismatches into ordinary engineering requirements.

D. FRAGMENTATION BETWEEN THE TWO COMMUNITIES

The fourth weakness is fragmentation. The post-quantum and verifiable-computing communities cite each other sparingly: lattice SNARK papers borrow module-lattice parameterisations without inheriting the implementation hygiene of the post-quantum literature, and migration roadmaps omit integrity services whose pairing-based proofs will need replacement in the same migration window. No surveyed work delivers an end-to-end prototype in which channel, identity, storage, computation and audit are simultaneously quantum-safe and publicly verifiable; the closest systems stop at two of the five layers. Benchmark suites likewise diverge, so cross-paper comparison required normalisation throughout this review. Fragmentation also hides compounding costs: a hybrid handshake plus a SNARK proof plus MAC tags can multiply message mass several-fold, and only an integrated treatment would expose the compound penalty. The agenda implied by these four weaknesses - realistic benchmarks, implementation-grade security, named and composed trust models, and integrated stacks - structures the discussion that follows. The cost of fragmentation is visible in this paper itself: assembling the reference architecture required drawing on five disjoint literatures, and no citation chain yet connects the FIPS standards to the verifiable fully homomorphic encryption constructions they will eventually have to protect. Until shared benchmarks exist, every cross-system claim - including this review's - carries that caveat.

5. DISCUSSION

The synthesis supports three claims. First, the components of a quantum-safe verifiable cloud already exist: standardised KEMs and signatures for channels and identities, transparent or lattice-based proof systems for computation, homomorphic MACs and authenticated structures for storage, and verifiable fully homomorphic encryption for the confidentiality-preserving frontier. Second, no single family suffices; the SIKE collapse and the still-young cryptanalysis of module lattices argue for portfolio diversity, which in turn demands crypto-agility as a first-class architectural property - algorithm registries, negotiable cipher suites and versioned proof-system identifiers, exactly the machinery sketched in Figure 2. Third, the binding constraint is not peak

performance but composition: every layer adds artefacts, and the sum must still fit the latency and bandwidth budgets of real tenants, which is a system co-design problem that no single cryptography paper can solve. Measurement is the practical starting point: operators cannot migrate what they have not inventoried, and the cryptographic bill of materials must be extended to verification dependencies - proof systems, commitment schemes and attestation roots - before any deadline becomes actionable. The roadmap of Figure 5 makes the sequencing explicit: hybrid transport now, composite PKI next, quantum-safe verification throughout, and convergence before the compliance window closes. The practical implications are concrete. Operators should inventory cryptographic and verification dependencies together, prioritise long-confidentiality data against harvest-now-decrypt-later, and adopt hybrid transport now while treating pairing-based integrity services as technical debt. Researchers should publish prover costs, side-channel evaluations and negotiation state-machine tests as first-class artefacts, standardise verifiable-computing benchmarks at common security levels, and pursue lattice- and hash-based proof systems whose assumptions inherit the scrutiny applied to FIPS candidates.

Regulators, finally, should recognise integrity as part of quantum readiness: a cloud that encrypts quantum-safely but verifies classically is only half migrated. Economics deserves equal attention: the surveyed cost models imply tiered verifiable services cheap homomorphic-MAC integrity for storage, succinct proofs for high-value batch analytics, verifiable fully homomorphic encryption reserved for regulated confidentiality-critical workloads and pricing these tiers against breach risk rather than raw CPU is an open product and research question. Standardisation of verifiable-computing interfaces - proof formats, verification-key distribution and attestation chains - should begin now, mirroring the early NIST process that made post-quantum cryptography deployable across vendors. Mosca's inequality makes the sequencing explicit - wherever confidentiality lifetime plus migration time approaches the plausible arrival of a cryptographically relevant quantum computer, migration of the verification stack must begin within the same programme, and the roadmap of Figure 5 shows that the window for doing so is open now but narrowing. Limitations of this review should also be acknowledged: the corpus is weighted towards lattice cryptography, hardware-attestation and blockchain-anchoring literatures appear only where they intersect cloud integrity, and part of the 2023 material is included at preprint quality, awaiting peer review. Taken together, the evidence reframes the question from whether the cloud can be verified to whether it can be verified in a way that survives the quantum era at acceptable cost.

6. CONCLUSION

This review examined past work on post-quantum cryptography and verifiable computing for cloud systems as a single, interdependent research programme. The standardisation of ML-KEM, ML-DSA and SLH-DSA, together with the draft FN-DSA and the selection of HQC, supplies quantum-safe primitives for channels, identities and storage; the verifiable computing lineage, from interactive proofs through succinct arguments and homomorphic MACs to verifiable fully homomorphic encryption, supplies the machinery by which a weak client can audit a powerful cloud. Our analysis shows that the two strands are technically mature in isolation but under-integrated in practice: evaluation bias, implementation leakage, unnamed trust assumptions and the absence of end-to-end quantum-safe verifiable prototypes are the recurring deficits. The path forward is a layered, agility-first architecture in which quantum-safe transport and signatures carry lattice- or hash-based

proofs, homomorphic authentication guards storage, and migration is governed by Mosca's theorem rather than by optimism. Until such stacks are benchmarked end to end, the cloud's two fundamental promises - that data stays secret and that results are true - remain only separately defensible. We contributed a unified taxonomy, a reference architecture, a protocol template, quantitative artefact comparisons and a roadmap anchored in thirty verified references, and the research agenda distilled here - integrated prototypes, implementation-grade evaluation and composed trust models - defines the work needed to defend them jointly. The central recommendation is simple: treat confidentiality and correctness as one migration, budgeted by Mosca's theorem and delivered by agile, layered, quantum-safe design.

7. REFERENCES

- [1] W. Diffie and M. Hellman, "New directions in cryptography," *IEEE Trans. Inf. Theory*, vol. 22, no. 6, pp. 644-654, Nov. 1976.
- [2] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proc. 35th Annu. Symp. Foundations of Computer Science (FOCS)*, Santa Fe, NM, USA, 1994, pp. 124-134.
- [3] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proc. 28th Annu. ACM Symp. Theory of Computing (STOC)*, Philadelphia, PA, USA, 1996, pp. 212-219.
- [4] M. Mosca, "Cybersecurity in an era with quantum computers: Will we be ready?" *IEEE Security Privacy*, vol. 16, no. 5, pp. 38-41, Sep./Oct. 2018.
- [5] D. J. Bernstein, J. Buchmann, and E. Buchmann, Eds., *Post-Quantum Cryptography*. Berlin, Germany: Springer-Verlag, 2009.
- [6] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," *J. ACM*, vol. 56, no. 6, pp. 1-40, 2009.
- [7] G. Alagic et al., "Status report on the third round of the NIST post-quantum cryptography standardization process," NIST, Gaithersburg, MD, USA, NIST IR 8413, 2022.
- [8] National Institute of Standards and Technology, "Module-Lattice-Based Key-Encapsulation Mechanism Standard," FIPS 203, Aug. 2023.
- [9] National Institute of Standards and Technology, "Module-Lattice-Based Digital Signature Standard," FIPS 204, Aug. 2023; "Stateless Hash-Based Digital Signature Standard," FIPS 205, Aug. 2023.
- [10] S. Goldwasser, Y. T. Kalai, and G. N. Rothblum, "Delegating computation: Interactive proofs for muggles," *J. ACM*, vol. 62, no. 4, pp. 1-64, 2015.
- [11] R. Gennaro, C. Gentry, and B. Parno, "Non-interactive verifiable computing: Outsourcing computation to arbitrary workers," in *Proc. CRYPTO*, Santa Barbara, CA, USA, 2010, pp. 465-482.
- [12] R. Gennaro, C. Gentry, B. Parno, and M. Raykova, "Quadratic span programs and succinct NIZKs without PCPs," in *Proc. EUROCRYPT*, Athens, Greece, 2013, pp. 626-645.

- [13] B. Parno, J. Howell, C. Gentry, and M. Raykova, "Pinocchio: Nearly practical verifiable computation," in Proc. IEEE Symp. Security and Privacy, Berkeley, CA, USA, 2013, pp. 238-252.
- [14] J. Groth, "On the size of pairing-based non-interactive arguments," in Proc. EUROCRYPT, Vienna, Austria, 2016, pp. 305-326.
- [15] R. Gennaro and D. Wichs, "Fully homomorphic message authenticators," in Proc. ASIACRYPT, Bengaluru, India, 2013, pp. 301-320.
- [16] C. Papamanthou, E. Shi, and R. Tamassia, "Verifiable delegation of computation with constant space and linear verifier time," in Proc. SCN, Amalfi, Italy, 2012, pp. 402-418.
- [17] M. Blum, W. S. Evans, P. Gemmell, S. Kannan, and M. Naor, "Checking the correctness of memories," *Algorithmica*, vol. 12, no. 2-3, pp. 225-244, 1994.
- [18] R. C. Merkle, "A digital signature based on a conventional encryption function," in Proc. CRYPTO, Santa Barbara, CA, USA, 1987, pp. 369-378.
- [19] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," *SIAM J. Comput.*, vol. 18, no. 1, pp. 186-208, 1989.
- [20] M. R. Albrecht, V. Cini, R. W. F. Lai, G. Malavolta, and S. A. Thyagarajan, "Lattice-based SNARKs: Publicly verifiable, preprocessing, and recursively composable," in Proc. 42nd Annu. Int. Cryptology Conf. (CRYPTO), 2022.
- [21] Y. Ishai, H. Su, and D. J. Wu, "Shorter and faster post-quantum designated-verifier zkSNARKs from lattices," in Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS), 2021, pp. 212-234.
- [22] E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev, "Scalable, transparent, and post-quantum secure computational integrity," *Cryptology ePrint Archive*, Rep. 2018/046, 2018.
- [23] D. F. Aranha, A. Costache, A. Guimaraes, and E. Soria-Vazquez, "HELIOPOLIS: Verifiable computation over homomorphically encrypted data from interactive oracle proofs is practical," in Proc. ACISP, 2023, pp. 1-24.
- [24] S. Atapoor, K. Baghery, H. V. L. Pereira, and J. Spiessens, "Verifiable FHE via lattice-based SNARKs," *IACR Commun. Cryptol.*, vol. 1, no. 1, 2023, doi: 10.62056/a6ksdkp10.
- [25] M. Chou and P. Cao, "Network impact of post-quantum certificate chain sizes on time to first byte in TLS deployments," *arXiv:2604.24869*, Apr. 2023.
- [26] P. Schwabe, D. Stebila, and T. Wiggers, "Post-quantum TLS without handshake signatures," in Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS), 2020, pp. 1445-1460.

- [27] A. A. Fall, "SoK: Systematizing hybrid strategies for the transition to post-quantum cryptography," IACR Cryptol. ePrint Arch., Rep. 2023/2052, 2023.
- [28] National Security Agency, "Quantum computing and post-quantum cryptography: CNSA Suite 2.0," Information Advisory, Sep. 2022, updated 2023.
- [29] Open Quantum Safe Project, "liboqs: An open-source library for post-quantum cryptographic algorithms," 2023. [Online]. Available: <https://github.com/open-quantum-safe/liboqs>
- [30] Z. Liu, A. Wang, C. Wei, Y. Ding, J. Zhang, A. Liu, and L. Zhu, "Release the power of rejected signatures: An efficient side-channel attack on the ML-DSA cryptosystem," IACR Cryptol. ePrint Arch., Rep. 2023/582, 2023.