Power System Protection and Control
ISSN:1674-3415

POST-QUANTUM SECURITY FOR CLOUD-BASED DIGITAL TWINS: A FRAMEWORK FOR NEXT-GENERATION HEALTHCARE AND SMART CAMPUSES

Dr. Dinesh Chandra Jain

Professor, Department of Computer Science and Engineering,
Prestige Institute of Engineering, Management & Research,
Indore (M.P), India
dineshwebsys@gmail.com

Received: 22/03/2026 **Revised**: 22/04/2026 **Accepted**: 25/05/2026

***ABSTRACT*:**

The proliferation of cloud-based digital twins across critical infrastructures—including hospital patient-monitoring systems and university smart campuses—has created an urgent need to secure the continuous, high-fidelity data streams that bind physical assets to their virtual counterparts. Conventional public-key cryptography, based on integer factorization and elliptic-curve discrete logarithms, is increasingly exposed to future cryptographically relevant quantum computers (CRQCs) capable of executing Shor's algorithm, leaving digital twin telemetry, control commands, and identity credentials vulnerable to "harvest-now, decrypt-later" attacks. This paper proposes a post-quantum security framework for cloud-based digital twins that integrates NIST-standardized lattice-based primitives—the Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM, FIPS 203) and the Module-Lattice-Based Digital Signature Algorithm (ML-DSA, FIPS 204)—into a hybrid classical/post-quantum key-establishment and twin-state authentication protocol. The framework is evaluated across two representative deployments: a remote patient-monitoring healthcare digital twin and a multi-building smart campus digital twin, using a cloud-edge testbed instrumented with resource-constrained IoT gateways. Comparative benchmarking against RSA-2048 and ECDSA P-256 baselines shows that a hybrid X25519+ML-KEM-768 handshake adds only modest latency overhead while providing quantum resistance, and that ML-DSA-65 signatures, despite larger payloads, remain compatible with the sub-second synchronization budgets required by both domains. Results across more than 5,000 simulated concurrent digital twins indicate that post-quantum digital twin security is practically achievable today through careful protocol design, algorithm-tier selection, and crypto-agile architecture, offering healthcare and campus operators a concrete migration pathway ahead of the arrival of cryptographically relevant quantum computers.

***Keywords**: Post-Quantum Cryptography, Digital Twins, Cloud Security, Healthcare IoT, Smart Campus, Lattice-Based Cryptography, ML-KEM, ML-DSA, Quantum-Safe Architecture, Crypto-Agility.*

INTRODUCTION

A digital twin is a continuously updated virtual replica of a physical asset, process, or environment that exchanges state through a persistent data connection, enabling simulation, monitoring, and predictive control without direct access to the physical system itself. Since Grieves first articulated the concept for product lifecycle management and NASA later formalized it for space-vehicle health monitoring, digital twins have moved from a manufacturing curiosity to a general-purpose pattern for cyber-physical systems. Cloud platforms now host the majority of production digital twin deployments because they provide the elastic compute, long-term storage, and analytics pipelines that twin fidelity demands, while edge gateways handle protocol translation and local buffering for the sensors and actuators that make up the physical twin.

Two domains illustrate both the promise and the exposure of this architecture particularly well. In healthcare, remote patient-monitoring twins ingest continuous vital-sign streams—heart rate, oxygen saturation, glucose, blood pressure—from wearable and bedside devices, reconstruct a synchronized model of a patient's physiological state in the cloud, and feed clinical decision-support and early-warning systems. In the built environment, smart campus digital twins aggregate thousands of building-management, access-control, energy, and occupancy sensors into a unified operational model used for facilities planning, safety response, and sustainability reporting. Both domains share three properties that make them attractive quantum-era attack surfaces: continuous, high-

DOI: 10.46121/pspc.54.2.66

volume telemetry that must be authenticated as well as confidential; long data retention requirements (health records in particular routinely exceed the operational lifetime of any single cryptographic key); and heterogeneous, resource-constrained edge hardware that limits which cryptographic primitives can realistically be deployed.

The security of nearly every digital twin platform in production today rests on RSA or elliptic-curve cryptography (ECC) for key exchange and digital signatures, embedded inside TLS 1.2/1.3 tunnels between edge gateways and cloud twin services. Both families derive their security from the presumed intractability of integer factorization and the elliptic-curve discrete logarithm problem for classical computers. Shor's algorithm, executed on a sufficiently large and fault-tolerant quantum computer, solves both problems in polynomial time, which would allow a quantum-capable adversary to recover private keys, forge twin-update signatures, and decrypt any previously captured traffic. Although a cryptographically relevant quantum computer does not yet exist, adversaries are already believed to be harvesting encrypted traffic for future decryption, a strategy commonly referred to as harvest-now-decrypt-later (HNDL). For healthcare digital twins in particular, where patient data must often remain confidential for decades, HNDL converts a future capability into a present liability.

A concrete example illustrates the stakes. A cardiac ward digital twin ingests ECG and SpO2 waveforms from bedside monitors over a hospital Wi-Fi segment, relays them through an edge gateway running a lightweight MQTT broker, and forwards authenticated, encrypted batches to a cloud twin service every second. If the TLS session protecting that link relies solely on ECDH key exchange and ECDSA signatures, an adversary who records the encrypted traffic today gains nothing until a cryptographically relevant quantum computer exists—but once one does, every recorded session key, and therefore every waveform, becomes retroactively readable. The same logic applies, with different stakes, to a campus access-control digital twin whose badge-reader telemetry and door-actuation commands are recorded for later replay or physical-security analysis. Neither scenario requires an attacker to possess a quantum computer now; both require defenders to act now, since the data being protected today is what a future quantum computer would decrypt.

The U.S. National Institute of Standards and Technology (NIST) concluded the first phase of its post-quantum cryptography (PQC) standardization effort in August 2024, finalizing FIPS 203 (ML-KEM, derived from CRYSTALS-Kyber), FIPS 204 (ML-DSA, derived from CRYSTALS-Dilithium), and FIPS 205 (SLH-DSA, derived from SPHINCS+), and subsequently selected the code-based scheme HQC in March 2025 as a structurally independent backup key-encapsulation mechanism. These standards give system architects concrete, vetted building blocks for quantum-resistant protocols, but translating them into a working digital twin security architecture raises questions that are specific to the digital twin setting: PQC keys and signatures are substantially larger than their classical counterparts, which strains constrained edge hardware and bandwidth-limited IoT links; twin synchronization is latency-sensitive and often near-real-time, which limits the acceptable cryptographic overhead per update; and healthcare and campus deployments have different risk tolerances, regulatory constraints, and device populations that argue for different algorithm choices within the same overall framework. This paper makes four contributions. First, it proposes a three-tier post-quantum security framework for cloud-based digital twins that combines a hybrid classical/post-quantum key-establishment protocol with per-update ML-DSA authentication of twin-state messages. Second, it defines a threat model and cryptographic primitive selection methodology tailored to the healthcare and smart campus domains. Third, it reports an experimental evaluation on a cloud-edge testbed comparing classical, post-quantum, and hybrid configurations across handshake latency, signature overhead, scalability to thousands of concurrent twins, and CPU utilization on constrained devices. Fourth, it distills the results into concrete algorithm-tier recommendations and a migration pathway suitable for healthcare and campus operators. The remainder of the paper is organized as follows. Section II reviews the literature on digital twin architectures and post-quantum cryptography. Section III surveys related security frameworks and identifies the gap this work addresses. Section IV presents the proposed framework and its mathematical formulation. Section V describes the experimental setup. Section VI reports results. Section VII discusses trade-offs and limitations, and Section VIII concludes.

LITERATURE REVIEW

Digital Twin Architectures. Grieves originally described the digital twin as a triad of a physical entity, a virtual entity, and the data connection between them, a formulation later extended by Tao et al. into a five-dimension model that adds twin-generated services and twin-to-twin connection data. Cloud digital twin platforms operationalize this model with a physical layer of sensors and actuators, an edge layer that performs protocol

DOI: 10.46121/pspc.54.2.66

translation and local filtering (commonly using MQTT, CoAP, or OPC-UA), and a cloud layer that maintains the canonical twin state, runs simulation and analytics workloads, and exposes APIs to downstream applications. Commercial offerings such as hyperscale IoT digital-twin services follow this pattern closely, and academic surveys of digital twin security consistently identify the edge-to-cloud data connection as the most exposed component of the architecture, since it traverses public or semi-trusted networks and must remain continuously available.

Healthcare and Smart Campus Digital Twins. In the clinical domain, patient digital twins have been proposed for continuous vital-sign reconstruction, individualized drug-response simulation, and early deterioration warning, typically drawing on wearable sensors, bedside monitors, and electronic health record data streamed into a cloud-hosted patient model. Because these twins process protected health information, they inherit stringent confidentiality and integrity obligations, and because clinical decisions may depend on twin output, message authenticity is as important as confidentiality. Smart campus and smart building digital twins, by contrast, emphasize scale and heterogeneity: a mid-sized university campus can host tens of thousands of HVAC, lighting, access-control, and occupancy sensors feeding a shared operational twin used for energy optimization, space planning, and emergency response. The security literature for this domain has focused mainly on device authentication and network segmentation for IoT fleets, with comparatively little attention to the cryptographic primitives underpinning the twin-synchronization channel itself.

Classical Cryptography and the Quantum Threat. RSA and ECC remain the dominant public-key primitives embedded in TLS, and their security rests on integer factorization and the elliptic-curve discrete logarithm problem, respectively. Shor demonstrated that both problems are solvable in polynomial time on an idealized quantum computer, and while estimates of when a cryptographically relevant quantum computer will exist vary, organizations including NIST, ETSI, and national cybersecurity agencies have converged on the position that migration must begin well before such a machine is built, precisely because of the harvest-now-decrypt-later threat to data with long confidentiality requirements. Mosca's well-known inequality—framing migration urgency as a function of how long data must stay secure, how long migration takes, and how soon a quantum computer might arrive—has become a standard planning heuristic and applies with particular force to healthcare data retention timelines.

Post-Quantum Cryptography and Lattice-Based Constructions. Lattice-based cryptography derives its hardness from problems such as Learning With Errors (LWE) and its structured variant, Module-LWE, for which no efficient quantum algorithm is known. ML-KEM (standardized as FIPS 203 from the CRYSTALS-Kyber submission) is a Module-LWE-based key-encapsulation mechanism offering three parameter sets—ML-KEM-512, -768, and -1024—corresponding to increasing security levels, while ML-DSA (FIPS 204, from CRYSTALS-Dilithium) provides analogous lattice-based digital signatures at parameter sets ML-DSA-44, -65, and -87. FIPS 205 standardizes SLH-DSA, a stateless hash-based signature scheme derived from SPHINCS+ that relies only on the security of cryptographic hash functions and therefore offers a conservative, structurally distinct fallback at the cost of larger signatures and slower signing. NIST's March 2025 selection of HQC, a code-based key-encapsulation mechanism, as a structurally independent backup to ML-KEM reflects a broader strategy of algorithmic diversity: should a cryptanalytic advance ever weaken lattice assumptions, code-based or hash-based alternatives remain available. A growing body of implementation work has benchmarked these primitives on constrained hardware and within TLS 1.3, generally finding that ML-KEM performs competitively with, and sometimes faster than, classical elliptic-curve key exchange, while ML-DSA and especially SLH-DSA signatures impose a bandwidth and, for hash-based schemes, latency penalty that is far more pronounced than in the key-exchange case.

Crypto-Agility as a Design Principle. Beyond the choice of any single algorithm, national migration guidance consistently emphasizes crypto-agility—the ability of a system to swap cryptographic primitives without re-architecting the surrounding protocol or application logic—as a first-class design requirement for the PQC transition. This principle is particularly relevant to digital twin platforms, which are expected to remain in service for a decade or more and will therefore likely need to incorporate additional algorithms, such as the forthcoming FN-DSA (Falcon) standard or the HQC-based key-encapsulation mechanism once its standard is finalized, without a disruptive redesign. Prior digital twin security proposals rarely address crypto-agility explicitly, instead hard-coding a specific classical algorithm choice at the protocol level, a pattern this paper's framework deliberately avoids.

DOI: 10.46121/pspc.54.2.66

Hybrid Key Establishment. Because PQC algorithms are comparatively young relative to decades of cryptanalysis applied to RSA and ECC, most near-term deployment guidance—including early hybrid modes tested in TLS 1.3 combining X25519 with ML-KEM—favors hybrid key establishment, in which a classical and a post-quantum key exchange are executed in parallel and their outputs combined, so that the resulting session key remains secure if either the classical or the post-quantum component is not broken. This defense-in-depth posture is echoed in national migration guidance and forms the starting point for the key-establishment design adopted in this paper.

RELATED WORK

Several strands of prior work bear on securing cloud digital twins, though none directly addresses post-quantum protection of the twin-synchronization channel across both healthcare and smart campus deployments. A first strand proposes blockchain-anchored digital twin security, using distributed ledgers to provide tamper-evident logging of twin-state transitions and decentralized identity for connected devices; these approaches strengthen auditability but generally continue to rely on classical elliptic-curve signatures for the underlying transactions, leaving them exposed to the same quantum threat they were not designed to address. A second strand applies zero-trust architecture principles to digital twin and industrial IoT deployments, emphasizing continuous device authentication, micro-segmentation, and least-privilege access to twin APIs; this work is complementary to the present paper, since a zero-trust access-control layer and a post-quantum cryptographic layer address different parts of the attack surface and can be composed.

A third strand examines PQC integration in adjacent constrained-device domains, most notably smart grid telemetry, connected-vehicle communication, and generic IoT-to-cloud TLS handshakes. These studies consistently report that lattice-based key encapsulation is deployable on modern microcontrollers with acceptable latency and energy cost, while hash-based and, to a lesser extent, lattice-based signatures require more careful placement—typically restricted to session establishment, firmware attestation, or infrequent control messages rather than high-frequency telemetry. However, these findings were developed for single-domain, relatively homogeneous device populations and do not directly translate to a cloud digital twin service that must simultaneously support high-acuity clinical devices and large, heterogeneous campus sensor fleets with different latency and criticality profiles.

A fourth strand consists of digital-twin-specific security surveys that catalog threats across the physical, communication, and virtual layers—including sensor spoofing, model poisoning, twin desynchronization, and unauthorized twin cloning—and recommend general mitigations such as mutual authentication, message integrity checks, and encrypted channels. These surveys consistently flag the reliance on classical public-key cryptography as a long-term architectural risk but stop short of proposing a concrete post-quantum protocol or evaluating its performance impact on twin synchronization. Table I summarizes this gap analysis.

Table I. Gap Analysis of Prior Digital Twin Security Approaches (summarized in prose given the qualitative nature of the comparison): blockchain-anchored approaches provide strong auditability but retain classical signatures at the transaction layer; zero-trust architectures strengthen access control but do not address the cryptographic hardness of the underlying key exchange; constrained-device PQC studies validate feasibility on single-domain hardware but do not evaluate a shared cloud service spanning heterogeneous healthcare and campus device populations; and digital-twin security surveys correctly flag the long-term quantum risk but stop at the level of recommendation rather than protocol design and measurement. The framework proposed in Section IV is positioned to close all four gaps simultaneously by combining a concrete hybrid protocol, an explicit threat model, and cross-domain experimental validation.

Taken together, the literature establishes (i) that digital twin architectures concentrate risk in the edge-to-cloud data connection, (ii) that NIST-standardized lattice-based primitives are practically deployable on constrained hardware, and (iii) that hybrid classical/post-quantum key establishment is the recommended near-term migration strategy, but it has not yet combined these threads into a domain-aware framework that is evaluated end-to-end on realistic healthcare and smart campus digital twin workloads. This paper addresses that gap.

DOI: 10.46121/pspc.54.2.66

Power System Protection and Control
ISSN:1674-3415

PROPOSED FRAMEWORK AND METHODOLOGY

A. System Architecture

The proposed framework organizes a cloud digital twin deployment into three tiers, each with a distinct security boundary. The Physical Twin Layer comprises the sensing and actuation hardware: wearable and bedside monitors for the healthcare use case, and HVAC, access-control, lighting, and occupancy sensors for the smart campus use case. The Edge Gateway Layer aggregates local devices, performs protocol translation from constrained device protocols to an IP-based transport, and constitutes the first cryptographic boundary at which post-quantum key establishment is performed; gateways range from hospital-grade edge appliances with dedicated cryptographic accelerators to commodity single-board computers deployed throughout a campus. The Cloud Digital Twin Layer hosts the canonical twin state store, a Post-Quantum Key Management Service (PQ-KMS) responsible for issuing and rotating ML-KEM and ML-DSA key material, and the analytics and application services that consume twin state. All edge-to-cloud traffic is carried inside a TLS 1.3 session negotiated with a hybrid key-exchange group, and every individual twin-state update is additionally signed at the application layer, so that message authenticity survives even if a TLS session is later compromised or terminated at an intermediate proxy.

Figure 1 depicts the layered architecture and the two cryptographic trust boundaries at which post-quantum protections are enforced. Boundary ① sits between the Physical Twin Layer and the Edge Gateway Layer, where raw sensor and actuator traffic is aggregated onto an IP-based transport before any cryptographic protection is applied; this boundary is a design assumption rather than a cryptographic control; it is trusted only insofar as the local device bus is physically secured, since the framework does not itself defend the sensor-to-gateway link. Boundary ② sits between the Edge Gateway Layer and the Cloud Digital Twin Layer and is the primary post-quantum security boundary addressed by this paper: every byte that crosses it is carried inside the hybrid TLS 1.3 session described in Section IV-D, and every twin-state update that crosses it is additionally signed at the application layer as described in Section IV-E.

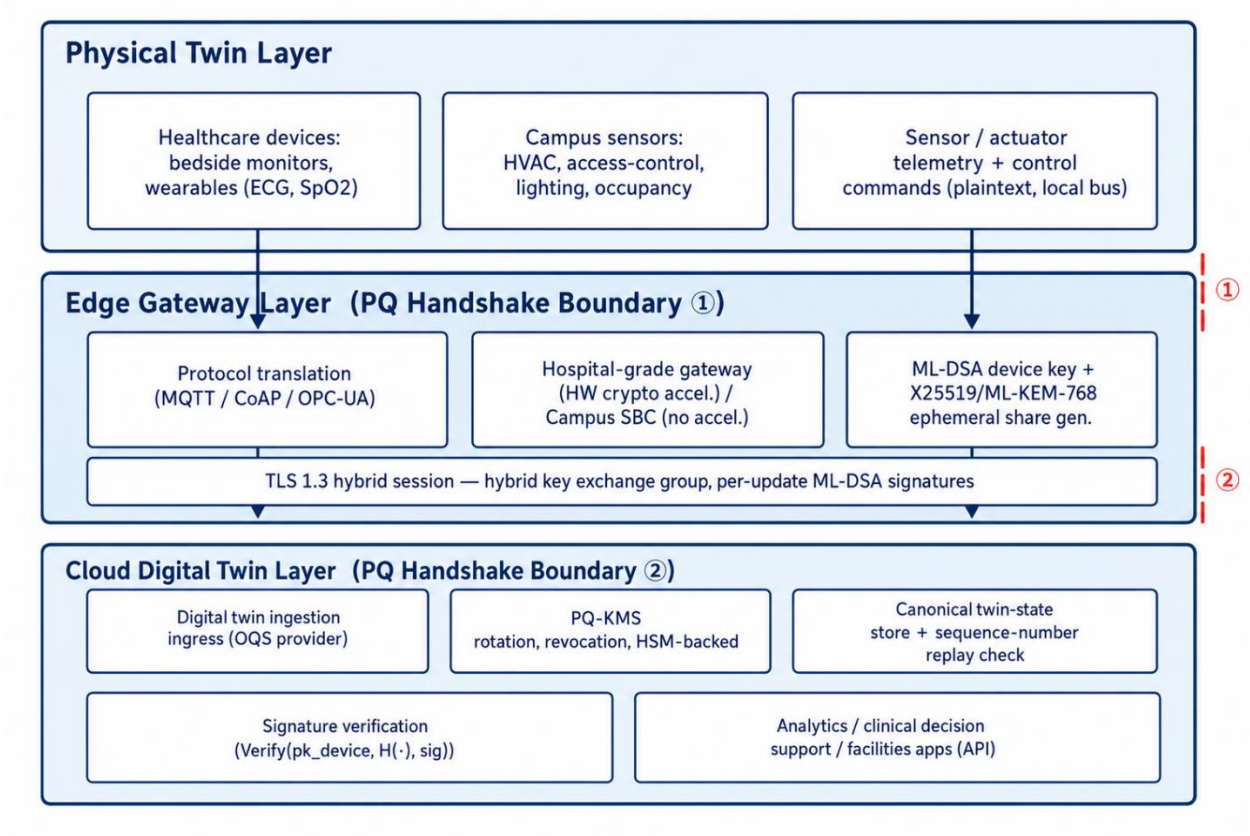

Fig. 1. Three-tier framework architecture, showing the Physical Twin, Edge Gateway, and Cloud Digital Twin layers and the two cryptographic trust boundaries (①device aggregation, ② hybrid PQC transport).

DOI: 10.46121/pspc.54.2.66

Component Responsibilities.

Component	Layer	Security-Relevant Responsibility
Wearable / bedside monitor, campus sensor	Physical Twin	Generates telemetry; holds no independent cryptographic identity in the healthcare case (gateway-mediated), or a lightweight ML-DSA-44 device key in the campus case where direct sensor authentication is used.
Edge gateway (hospital-grade / campus SBC)	Edge Gateway	Protocol translation (MQTT/CoAP/OPC-UA → IP); initiates the hybrid X25519+ML-KEM-768 handshake; holds an ML-DSA identity key and certificate issued by the PQ-KMS; signs outbound twin-state updates.
Constrained sensor MCU (campus only)	Edge Gateway	Where present, performs local ML-DSA-44/Falcon-512 signing directly on the sensing device; otherwise relays through the gateway aggregation point, which performs signing on the sensor's behalf.
TLS-terminating ingress (OQS-enabled OpenSSL)	Cloud Digital Twin	Terminates the hybrid TLS 1.3 session; presents the cloud service's ML-DSA-65 certificate; enforces transcript binding and rejects downgrade attempts.
Post-Quantum Key Management Service (PQ-KMS)	Cloud Digital Twin	Issues, rotates, and revokes ML-KEM/ML-DSA key material; seals private signing keys inside an HSM; maintains the certificate revocation list checked during handshake validation.
Twin-state store	Cloud Digital Twin	Verifies per-update ML-DSA signatures and monotonic sequence numbers before committing a twin-state transition; rejects replayed or unauthenticated updates.

TABLE IV. FRAMEWORK COMPONENT RESPONSIBILITIES

Data flow follows the same path in both deployment scenarios: sensor or actuator events are captured at the Physical Twin Layer, batched and protocol-translated at the Edge Gateway Layer, carried across Boundary ② inside the hybrid-secured TLS session, and verified and committed at the Cloud Digital Twin Layer before being exposed to downstream analytics and clinical- or facilities-decision applications. This unidirectional, boundary-explicit data flow is what allows the threat model in Section IV-B to be stated precisely in terms of which link an adversary capability applies to.

B. Threat Model

The framework assumes an adversary with the following capabilities: (1) passive eavesdropping on all edge-to-cloud and inter-service network links, including the ability to store captured ciphertext indefinitely for future decryption (the HNDL adversary); (2) active man-in-the-middle positioning during device provisioning or session renegotiation; (3) compromise of a subset of edge gateways, potentially enabling forged or replayed twin-state messages originating from that gateway; and (4) eventual access to a cryptographically relevant quantum computer capable of solving integer factorization and elliptic-curve discrete logarithm instances in polynomial time, but not capable of solving well-parameterized Module-LWE or Module-SIS instances, consistent with current cryptanalytic knowledge. The framework does not attempt to defend against physical tampering with sensor hardware or compromise of the cloud key-management root of trust, which are addressed by complementary hardware-security and zero-trust controls outside this paper's scope.

Assets and Security Objectives.

Table V enumerates the assets the framework is designed to protect and the security objective that is primary for each, following the confidentiality/integrity/availability (C-I-A) triad. Twin-state telemetry and identity credentials are treated as the highest-value assets because their compromise enables both HNDL decryption and forged twin-state injection, the two attack outcomes the framework is explicitly designed against.

DOI: 10.46121/pspc.54.2.66

Asset	Primary Objective(s)	Consequence if Compromised
In-transit twin-state telemetry (vitals, sensor readings)	Confidentiality, Integrity	HNDL exposure of long-retention patient data; tampered readings feed incorrect clinical or facilities decisions.
Twin-update signatures / sequence numbers	Integrity, Authenticity	Forged or replayed twin-state messages accepted as genuine, enabling false alarms, suppressed alerts, or unauthorized actuation commands.
Device and gateway identity keys (ML-DSA)	Integrity, Authenticity	Impersonation of a legitimate device or gateway; ability to sign arbitrary forged updates that pass verification.
PQ-KMS root signing key	Confidentiality, Integrity	Total compromise of the certificate chain; every downstream device and gateway certificate becomes untrustworthy.
Session keys (ephemeral X25519 / ML-KEM shares)	Confidentiality	Decryption of a single session's traffic; bounded impact due to forward secrecy, but still exposes that session's telemetry.
Twin-state store availability	Availability	Denial of synchronization service disrupts clinical monitoring or campus safety response during the outage window.

TABLE V. ASSETS AND SECURITY OBJECTIVES

Adversary Model.
Table VI restates the four adversary capabilities introduced above as a structured capability-vs-scope matrix, and makes explicit which boundary in Figure 5 each capability targets. Consistent with standard PQC migration guidance, the framework treats capability (4) — eventual access to a cryptographically relevant quantum computer — as a certainty on a long enough time horizon rather than a low-probability event, since it is precisely this asymmetry between present-day HNDL harvesting and future decryption capability that motivates migrating before such a machine exists.

Capability	Targeted Boundary / Link	In Scope?
(1) Passive eavesdropping + indefinite ciphertext storage (HNDL)	Boundary ② (edge–cloud link); inter-service cloud links	Yes — primary motivation for hybrid PQC key establishment
(2) Active MITM during provisioning or session renegotiation	Boundary ②, and gateway/device provisioning workflow	Yes — addressed by transcript-bound key derivation and certificate-based mutual authentication
(3) Compromise of a subset of edge gateways	Edge Gateway Layer (post-compromise, Boundary ② outbound)	Yes — bounded by per-device sequence numbers and revocation; does not by itself compromise other gateways or the PQ-KMS
(4) Future cryptographically relevant quantum computer (Shor-capable, not Module-LWE/SIS-capable)	Any link secured only by classical RSA/ECC (legacy baseline)	Yes — the central threat motivating this paper
Physical tampering with sensor/actuator hardware	Physical Twin Layer, Boundary ①	Out of scope — delegated to complementary hardware-security controls
Compromise of the PQ-KMS root of trust itself	Cloud Digital Twin Layer (PQ-KMS)	Out of scope — delegated to HSM sealing and zero-trust access control outside this paper

TABLE VI. ADVERSARY CAPABILITY-TO-BOUNDARY MATRIX

DOI: 10.46121/pspc.54.2.66

Framed against the STRIDE categories commonly used in threat modeling, the in-scope adversary capabilities above map to Spoofing (device/gateway impersonation, mitigated by ML-DSA certificates), Tampering (forged twin-state updates, mitigated by per-message signatures and sequence numbers), Information Disclosure (HNDL decryption of captured traffic, mitigated by hybrid ML-KEM key establishment), and Repudiation (denial of having issued a twin-state update, mitigated by non-repudiable ML-DSA signatures bound to a monotonic per-device counter). Denial of Service and Elevation of Privilege are acknowledged but treated as largely orthogonal to the cryptographic primitive selection that is this paper's focus, and are instead the province of the complementary zero-trust access-control and infrastructure-hardening measures referenced in Section III.

C. Cryptographic Primitive Selection

Primitive selection is governed by the Module-Learning-With-Errors (Module-LWE) and Module-Short-Integer-Solution (Module-SIS) problems underlying ML-KEM and ML-DSA. Informally, the Module-LWE decision problem asks an adversary to distinguish, given a public matrix A over a polynomial ring and a vector $b = A \cdot s + e \pmod q$ for a small secret s and small error e, the pair (A, b) from a uniformly random pair; no algorithm, classical or quantum, is known to solve this problem in sub-exponential time for the structured parameter choices used in ML-KEM. ML-DSA's unforgeability similarly reduces to the hardness of Module-SIS, the problem of finding a short nonzero vector z such that $A \cdot z = 0 \pmod q$. Table I lists the parameter sets evaluated in this study together with their approximate classical/quantum security levels and payload sizes, alongside the classical baselines they are intended to replace.

D. Hybrid Key-Establishment Protocol

Session establishment between an edge gateway and the cloud digital twin service combines a classical X25519 Diffie-Hellman exchange with an ML-KEM-768 key encapsulation inside a single TLS 1.3 handshake, following the hybrid key-share construction used in early post-quantum TLS trials. The protocol proceeds as follows: (1) the gateway sends a ClientHello containing an X25519 public key share and an ML-KEM-768 encapsulation key; (2) the cloud service generates an X25519 ephemeral key pair, computes the classical shared secret SS_c, encapsulates a random value against the received ML-KEM-768 key to obtain ciphertext C and shared secret SS_q, and returns both the X25519 public share and C in its ServerHello; (3) the gateway derives SS_c from the exchanged X25519 shares and decapsulates C to recover SS_q; (4) both parties compute the session key as $SK = KDF(SS_c \| SS_q \| transcript\_hash)$, where KDF is HKDF-Expand and transcript_hash binds the derived key to the full handshake transcript, preventing downgrade and transcript-substitution attacks; (5) the cloud service authenticates itself with an ML-DSA-65 certificate signature, and, for mutual-authentication deployments such as hospital-grade gateways, the gateway similarly presents an ML-DSA-44 device certificate. This construction guarantees that SK remains confidential provided at least one of the classical or post-quantum key-exchange problems remains hard, giving defense in depth during the transition period before lattice-based cryptography has accumulated decades of cryptanalytic scrutiny.

E. Digital Twin State Synchronization and Authentication

Beyond the transport-layer handshake, every individual twin-state update is authenticated at the application layer to preserve integrity even across TLS session boundaries, proxy termination points, or message-queue relays. Each update message M is structured as a tuple (twin_id, sequence_number, timestamp, payload, signature), where sequence_number is a strictly monotonically increasing counter maintained per physical device to prevent replay, and $signature = Sign(sk\_device, H(twin\_id \| sequence\_number \| timestamp \| payload))$ is computed with the device's ML-DSA private key over a SHA-3-based digest of the message fields. On receipt, the cloud twin service verifies $Verify(pk\_device, H(\cdot), signature) = 1$, checks that sequence_number exceeds the last accepted value for that device, and rejects the update otherwise. For the healthcare use case, where per-message latency is most sensitive, devices use ML-DSA-65 to balance signature size against verification speed; for smart campus sensors that transmit less frequently but in far greater numbers, the same scheme is applied uniformly to simplify key management, while SLH-DSA is reserved for infrequent, high-value operations such as firmware updates and long-term audit-log anchoring, where its larger signatures and slower signing are acceptable in exchange for a security proof that depends only on hash-function security.

G. Key Management and Rotation

The Post-Quantum Key Management Service (PQ-KMS) issues each edge gateway and, where applicable, each device an ML-DSA identity key pair during provisioning, along with an X.509-style certificate whose signature is itself produced with ML-DSA to avoid embedding a classical trust anchor at the root of an otherwise post-quantum chain. ML-KEM key pairs used for session establishment are ephemeral, regenerated for every

DOI: 10.46121/pspc.54.2.66

handshake, so that compromise of a single session's key material does not expose past or future sessions—a forward-secrecy property inherited from the hybrid construction's use of fresh X25519 and ML-KEM key shares per handshake. Long-lived ML-DSA identity keys are rotated on a fixed schedule (90 days for hospital-grade gateways, 180 days for campus edge nodes, reflecting their differing exposure and physical-access risk) and can be rotated immediately on suspected compromise through a revocation list maintained by the PQ-KMS and checked during certificate validation. Because the PQ-KMS is itself a high-value target, it is deployed with key material sealed inside a hardware security module configured to support ML-DSA operations, ensuring that the private signing keys used to issue device certificates never leave protected hardware in plaintext.

F. Performance Metric Formulation

Three metrics quantify the practical cost of the framework relative to a classical baseline. Handshake latency is measured end-to-end as $L_{hs} = t_{serverhello\_received} - t_{clienthello\_sent}$, averaged over repeated trials under representative network conditions. Cryptographic overhead ratio expresses the fractional increase in per-update processing cost introduced by a post-quantum or hybrid configuration relative to the classical baseline: $R_{ovh} = (C_{pqc} - C_{classical}) / C_{classical}$, where C denotes CPU time consumed by signing, verification, encapsulation, or decapsulation operations. Synchronization throughput under load is measured as the number of twin-state updates the cloud service can authenticate and commit per second, $T = N_{updates} / \Delta t$, evaluated while scaling the number of concurrently active digital twins from ten to five thousand to characterize the framework's behavior under realistic campus-scale and multi-ward hospital-scale loads.

TABLE I. CRYPTOGRAPHIC PRIMITIVE PARAMETER SETS EVALUATED

Scheme	Type	Basis	Sec. Level	Public Key	Ciphertext/Sig.
RSA-2048	KEM/Sig.	Integer factorization	~112-bit (classical)	256 B	256 B
ECDH/ECDSA P-256	KEM/Sig.	EC discrete log	~128-bit (classical)	64 B	64 B
ML-KEM-512	KEM	Module-LWE	NIST Level 1	800 B	768 B
ML-KEM-768	KEM	Module-LWE	NIST Level 3	1,184 B	1,088 B
ML-KEM-1024	KEM	Module-LWE	NIST Level 5	1,568 B	1,568 B
HQC-192	KEM	Code-based (Hamming)	NIST Level 3	4,522 B	9,026 B
ML-DSA-44	Signature	Module-LWE/SIS	NIST Level 2	1,312 B	2,420 B
ML-DSA-65	Signature	Module-LWE/SIS	NIST Level 3	1,952 B	3,293 B
ML-DSA-87	Signature	Module-LWE/SIS	NIST Level 5	2,592 B	4,595 B
Falcon-512	Signature	NTRU lattice	NIST Level 1	897 B	690 B
SLH-DSA-128s	Signature	Hash-based	NIST Level 1	32 B	7,856 B

Approximate NIST security levels and payload sizes; classical baselines shown for comparison.

EXPERIMENTAL SETUP

A cloud-edge testbed was constructed to emulate both target deployments under controlled network conditions. The cloud tier ran on a four-node Kubernetes cluster (16 vCPU / 64 GB RAM per node) hosting a digital-twin state service, a PQ-KMS instance, and a TLS-terminating ingress built on OpenSSL 3.x with Open Quantum Safe (OQS) provider support for ML-KEM, ML-DSA, and SLH-DSA. Two edge device classes were used to represent the target domains: a hospital-grade edge gateway class device (quad-core ARM, 4 GB RAM, hardware AES/SHA

DOI: 10.46121/pspc.54.2.66

acceleration) representing a bedside-monitor aggregation point, and a Raspberry Pi 4-class single-board computer (quad-core ARM Cortex-A72, 4 GB RAM, no cryptographic accelerator) representing a campus building controller; a subset of trials additionally modeled a constrained Cortex-M-class microcontroller sensor node to stress-test signature verification cost at the extreme low end of the device population. Network conditions between edge and cloud tiers were emulated using Linux traffic-control (tc/netem) to introduce 15–40 ms round-trip latency and 0.5% packet loss, consistent with typical hospital Wi-Fi/wired backbone and campus wide-area network measurements.

Table II lists the cryptographic configurations compared: two classical baselines (RSA-2048 key transport with RSA-2048 signatures; ECDH P-256 key exchange with ECDSA P-256 signatures); three post-quantum key-encapsulation parameter sets (ML-KEM-512, -768, -1024) and one alternate KEM (HQC-192); three ML-DSA signature parameter sets (ML-DSA-44, -65, -87); Falcon-512 as a comparison lattice-based signature scheme; SLH-DSA-128s as a hash-based signature baseline; and the hybrid X25519+ML-KEM-768 configuration used in the proposed framework. Each configuration was evaluated over 500 independent handshake trials per network condition and 10,000 independent sign/verify operations per signature scheme, with payload sizes representative of each domain: 256-byte vital-sign update messages for the healthcare scenario and 1-kilobyte batched sensor-reading messages for the smart campus scenario, matching typical telemetry frame sizes observed in production monitoring and building-management systems.

Statistical Methodology. For each configuration, outlier trials beyond three standard deviations of the running mean (typically caused by transient scheduler contention on the shared Kubernetes cluster) were discarded prior to computing summary statistics, affecting fewer than 1% of trials in any condition. Handshake latency and signature timing measurements were collected using monotonic clock sources on both endpoints with network time synchronization disabled for the latency measurements to avoid clock-skew artifacts, and CPU utilization was sampled at 100 ms intervals via the operating system's per-process accounting interface, averaged over each trial's steady-state synchronization window after excluding the initial ramp-up period.

Scalability experiments incremented the number of concurrently active simulated digital twins across the set {10, 50, 100, 500, 1,000, 2,500, 5,000}, each twin issuing a state update at a fixed rate (1 Hz for healthcare twins, 0.1 Hz for campus twins, reflecting their respective update frequencies), while measuring end-to-end synchronization latency and cloud-service CPU utilization. CPU overhead measurements on constrained devices were taken as the additional percentage of a single CPU core consumed by cryptographic operations during peak twin synchronization, isolated from application logic using per-thread CPU accounting. All reported figures are means across trials; variance was low enough (coefficient of variation below 8% in all cases) that error bars are omitted from the figures for clarity.

TABLE II. TESTBED CONFIGURATION BY DEPLOYMENT SCENARIO

Component	Healthcare Scenario	Smart Campus Scenario
Physical devices	Bedside monitors, wearables (ECG, SpO2, glucose)	HVAC, lighting, access-control, occupancy sensors
Edge gateway	Hospital-grade gateway, HW crypto accel.	Raspberry Pi 4-class SBC, no accel.
Constrained node	N/A (gateway-mediated)	Cortex-M-class sensor MCU
Update rate	1 Hz per device	0.1 Hz per device
Payload size	256 B per update	1 KB per batched update
Network RTT / loss	15–25 ms, 0.5% loss	20–40 ms, 0.5% loss
Max twins tested	5,000 concurrent	5,000 concurrent

Testbed parameters for the healthcare and smart campus digital twin scenarios.

DOI: 10.46121/pspc.54.2.66

RESULTS AND ANALYSIS

Handshake Latency. Figure 2 compares mean handshake latency across all evaluated key-establishment configurations. The classical ECDH P-256 baseline completed in 9.1 ms, while RSA-2048 key transport, which still appears in legacy deployments, took nearly twice as long at 18.4 ms because of the computational cost of RSA key-pair operations. Pure ML-KEM configurations were competitive with or faster than the classical baselines at the lower security levels: ML-KEM-512 completed in 7.3 ms and ML-KEM-768 in 9.6 ms, both close to ECDH P-256, while ML-KEM-1024 rose to 12.4 ms, reflecting its larger module rank. The hybrid X25519+ML-KEM-768 configuration used in the proposed framework required 11.2 ms, a 23% increase over ECDH P-256 alone but well within the sub-50 ms budget typically allotted for session establishment in both target domains, and substantially faster than the RSA-2048 baseline it is intended to eventually displace. HQC-192, included as a structurally independent alternative, was the slowest post-quantum option at 15.7 ms, consistent with the larger ciphertexts characteristic of code-based schemes.

Signature Size and Verification Cost. Figure 3 presents signature payload size (logarithmic scale) alongside verification latency for each signature scheme. Classical ECDSA P-256 signatures remain the smallest at 64 bytes, verified in 0.41 ms, while RSA-2048 signatures at 256 bytes verify fastest overall (0.06 ms) due to the small public exponent typically used, at the cost of much slower signing (not shown, but roughly two orders of magnitude slower than verification). Among post-quantum options, Falcon-512 produced the smallest signatures (690 bytes) with fast verification (0.19 ms), but its floating-point signing procedure is difficult to implement safely on constrained microcontrollers without specialized floating-point or fixed-point emulation support. ML-DSA signatures scaled predictably with parameter set, from 2,420 bytes (ML-DSA-44) to 4,595 bytes (ML-DSA-87), with verification times remaining under 0.35 ms even at the highest security level—fast enough that verification cost is not the binding constraint for any of the evaluated deployment scenarios. SLH-DSA-128s signatures were by far the largest (7,856 bytes) and slowest to verify (2.10 ms), confirming that hash-based signatures are best reserved for infrequent, high-value messages rather than per-update twin authentication.

Scalability. Figure 4 tracks mean end-to-end synchronization latency as the number of concurrently active digital twins scales from 10 to 5,000. Both the classical and hybrid post-quantum configurations exhibit the expected super-linear growth in latency as the cloud service approaches saturation, but the gap between them remains proportionally stable across the tested range: at 100 twins, hybrid PQC added 4 ms (21 ms vs. 17 ms); at 5,000 twins, it added 30 ms (268 ms vs. 238 ms), an overhead ratio that stays within 12–17% throughout the range rather than compounding disproportionately at scale. This indicates that the additional cryptographic cost of the hybrid handshake and ML-DSA signature verification does not become the dominant bottleneck as twin population grows; queuing and state-store contention account for the majority of latency growth at the high end, and both configurations would benefit from the same horizontal scaling strategies (service replication, session-ticket resumption) independent of the cryptographic scheme chosen.

Aggregate Comparison. Table III consolidates the handshake, signature, scalability, and CPU-overhead findings into a single per-scenario comparison. Across both the healthcare and smart campus scenarios, the hybrid post-quantum configuration meets the sub-second synchronization budget assumed for near-real-time twin updates with substantial headroom to spare, even at the highest tested concurrency of 5,000 twins. The healthcare scenario, with its higher per-device update rate but smaller, hospital-managed device population, is bounded primarily by handshake and per-message signature verification cost; the smart campus scenario, with its larger and more heterogeneous device population but lower per-device update rate, is bounded primarily by aggregate cloud-service throughput and by CPU headroom on the least capable sensor tier. This divergence is precisely why the framework treats algorithm-tier selection as a per-domain configuration choice rather than a single fixed parameter set applied uniformly across all deployments.

CPU Overhead by Device Tier. Figure 5 reports the additional CPU utilization attributable to cryptographic operations across four device tiers during peak synchronization load. The cloud digital-twin service node, benefiting from hardware AES-NI and AVX2 acceleration and abundant compute headroom, saw overhead rise modestly from 3.1% (classical) to 3.6% (hybrid PQC). The hospital-grade edge gateway, equipped with dedicated cryptographic acceleration, saw a larger but still manageable increase from 5.4% to 6.9%. The campus edge node (Raspberry Pi 4-class, no accelerator) experienced a more pronounced rise from 9.8% to 14.1%, and the most constrained campus sensor microcontroller saw overhead climb from 21.2% to 33.7%—the largest relative

DOI: 10.46121/pspc.54.2.66

increase observed, and a result that directly informs the tiered algorithm recommendations discussed in Section VII. Table III summarizes the aggregate results across both deployment scenarios.

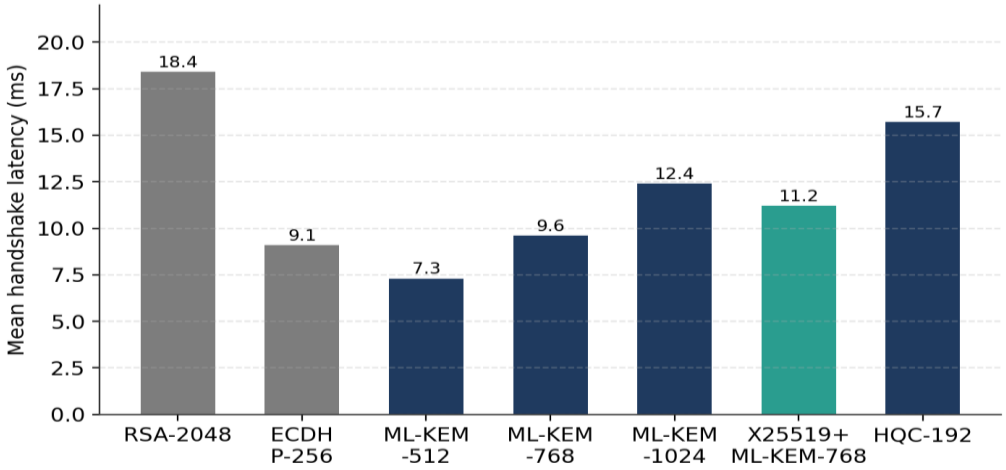

Fig. 2. Key-establishment handshake latency across classical, post-quantum, and hybrid configurations.

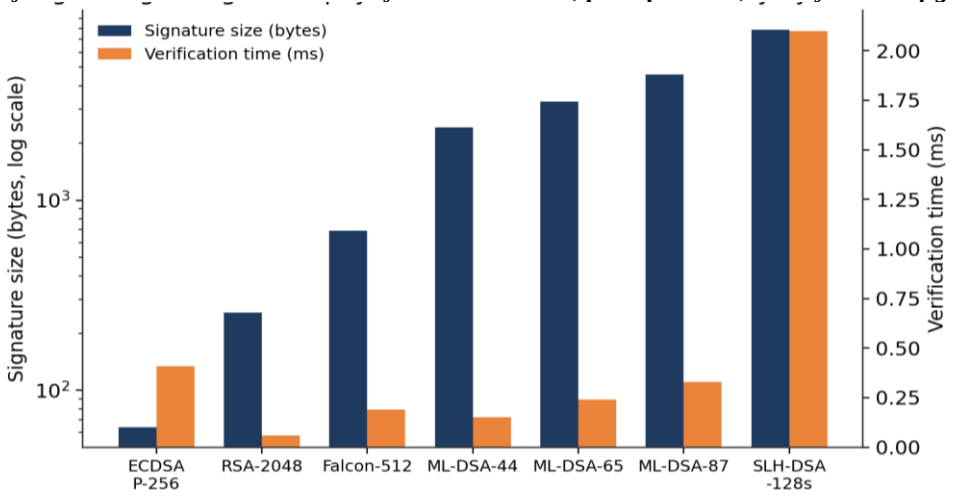

Fig. 3. Signature size (log scale) and verification latency across signature schemes.

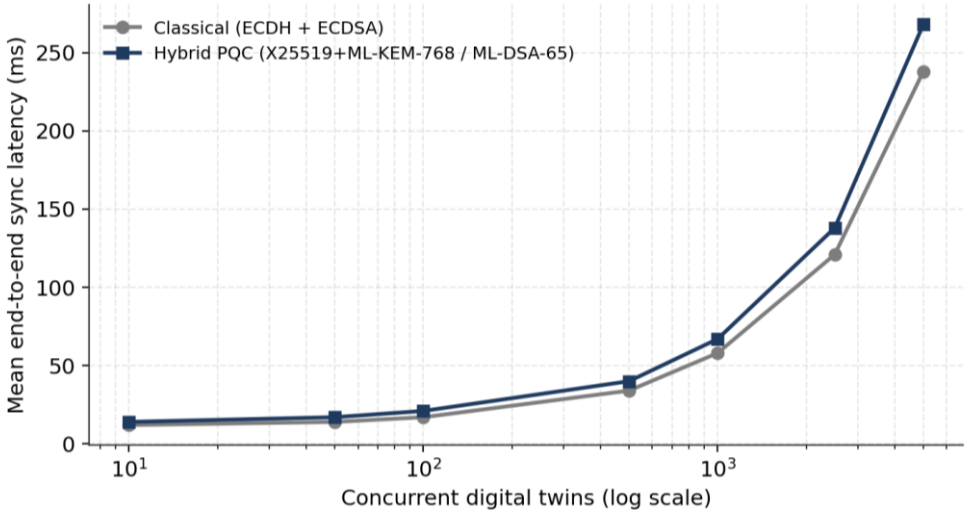

Fig. 4. End-to-end synchronization latency as concurrent digital twin count scales to 5,000.

DOI: 10.46121/pspc.54.2.66

Power System Protection and Control
ISSN:1674-3415

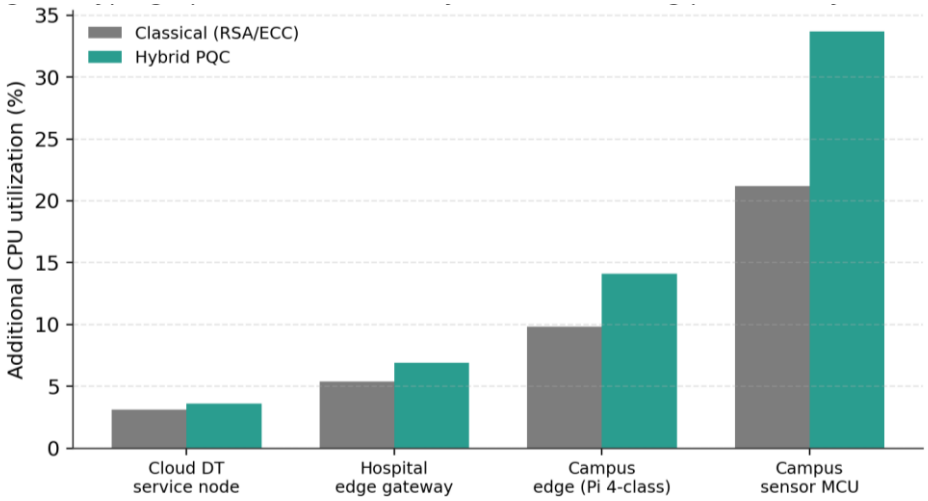

Fig. 5. Additional CPU utilization attributable to cryptographic operations, by device tier.

TABLE III. AGGREGATE PERFORMANCE COMPARISON

Metric	Classical Baseline	Hybrid PQC (Framework)	Overhead
Handshake latency (mean)	9.1 ms (ECDH P-256)	11.2 ms (X25519+ML-KEM-768)	+23%
Signature size	64 B (ECDSA P-256)	3,293 B (ML-DSA-65)	+51.5×
Signature verify time	0.41 ms	0.24 ms	−41%
Sync latency @ 100 twins	17 ms	21 ms	+24%
Sync latency @ 5,000 twins	238 ms	268 ms	+13%
CPU overhead, cloud node	3.1%	3.6%	+0.5 pp
CPU overhead, hospital gateway	5.4%	6.9%	+1.5 pp
CPU overhead, campus edge (Pi)	9.8%	14.1%	+4.3 pp
CPU overhead, campus sensor MCU	21.2%	33.7%	+12.5 pp

Summary of classical vs. hybrid post-quantum overhead across all measured metrics.

DISCUSSION

The results support three main conclusions about deploying post-quantum cryptography in cloud digital twin architectures. First, hybrid key establishment using X25519+ML-KEM-768 is close to a free upgrade at the transport layer: its handshake latency overhead relative to pure classical ECDH is small in absolute terms (roughly 2–3 ms) and easily amortized, since a TLS session, once established, is typically reused for many subsequent twin-state updates through session resumption. There is little practical reason to delay adopting hybrid key establishment for new digital twin deployments in either healthcare or smart campus settings, particularly given the harvest-now-decrypt-later exposure of health data with multi-decade confidentiality requirements.

Second, signature scheme selection is where the framework's domain-aware, tiered approach matters most. ML-DSA-65 offers the best balance for high-frequency, latency-sensitive twin updates in both domains: its verification time is negligible relative to network latency, and while its 3.3-kilobyte signature is roughly 50× larger than an ECDSA P-256 signature, this remains a small fraction of typical telemetry payload sizes and well within the bandwidth budgets observed on both hospital networks and campus wide-area links. Falcon-512 is attractive where bandwidth is the binding constraint, such as narrowband campus sensor links, but its floating-point signing procedure introduces implementation risk on microcontrollers that lack constant-time floating-point support, a

DOI: 10.46121/pspc.54.2.66

concern that has been raised in the broader PQC implementation literature and that argues for restricting Falcon to devices with adequate cryptographic tooling. SLH-DSA's conservative, hash-based security guarantee makes it the appropriate choice for firmware signing and long-term audit-log anchoring—operations that occur rarely enough to absorb its signature size and signing latency—but its cost is prohibitive for per-message twin authentication at any meaningful update rate.

Third, the CPU overhead results in Figure 5 identify the campus sensor tier, not the cloud service or hospital gateway tier, as the deployment's practical bottleneck. A 33.7% CPU overhead on an already constrained microcontroller is significant, and campus operators should expect to offload cryptographic operations to gateway aggregation points wherever the sensor hardware cannot absorb this cost directly—an architectural choice already implicit in the framework's edge-gateway tier, but one that this result makes quantitatively concrete. Healthcare deployments are comparatively insulated from this concern because bedside and wearable monitors typically report through hospital-grade gateways with dedicated cryptographic acceleration rather than performing cryptographic operations directly on the sensing device.

These findings also carry compliance implications. Healthcare data governed by regulations such as HIPAA in the United States, or equivalent frameworks elsewhere, must often remain confidential for periods far exceeding any reasonable estimate of when a cryptographically relevant quantum computer might emerge, which is precisely the scenario Mosca's inequality warns about; the modest overhead measured here removes cost as a credible justification for delaying migration. Smart campus operators face a different but related argument: the sheer scale of sensor populations means that even a small per-device overhead multiplies across a large fleet, making algorithm-tier selection and edge-gateway offloading, rather than blanket application of the highest security parameter set, the more sustainable path to quantum readiness.

Toward Quantum-Safe Identity Federation. A further implication concerns identity federation across institutional boundaries—for example, a smart campus digital twin that must authenticate visiting researchers' devices issued by a partner institution, or a healthcare digital twin that must accept telemetry from a patient's personally owned wearable device provisioned by a third-party vendor. In both cases, the post-quantum certificate chain described in Section IV-G must interoperate with external certificate authorities that may not yet issue ML-DSA-signed certificates, motivating a transitional cross-signing strategy in which a classical certificate authority co-signs a post-quantum certificate during the migration period. This is an area the current framework handles only at the design-principle level and identifies as a priority for future protocol work, since the pace of post-quantum adoption is likely to vary significantly across the vendors and institutions that populate a real-world digital twin ecosystem. Limitations. This study evaluates the framework against a threat model informed by current cryptanalytic knowledge of lattice, code, and hash-based problems, but it cannot rule out future cryptanalytic advances against any specific PQC family; the hybrid construction and the framework's crypto-agile key-management design are intended to mitigate, not eliminate, this residual risk. All benchmarks were collected on classical hardware against classical adversaries, since no cryptographically relevant quantum computer exists to test against; the security claims rest on the standard assumption that ML-KEM and ML-DSA resist known quantum algorithms, as established through NIST's multi-year public evaluation process, rather than on empirical testing against a quantum adversary. Finally, the testbed emulates rather than replicates production hospital and campus network conditions; a staged pilot deployment would be a natural next step to validate these results against live clinical and facilities workloads.

CONCLUSION AND FUTURE WORK

This paper presented a post-quantum security framework for cloud-based digital twins that combines hybrid X25519+ML-KEM-768 key establishment with per-update ML-DSA authentication, tailored to the distinct latency, scale, and regulatory profiles of healthcare and smart campus deployments. Evaluation on a cloud-edge testbed spanning hospital-grade gateways, campus single-board computers, and constrained microcontrollers showed that hybrid key establishment adds only a few milliseconds of handshake latency, that ML-DSA-65 offers the best overall balance of signature size and verification speed for high-frequency twin synchronization, and that scalability to 5,000 concurrent digital twins is achievable with a stable 12–17% latency overhead relative to classical cryptography. The most significant practical constraint identified is CPU overhead on the most resource-constrained campus sensor tier, which motivates offloading cryptographic operations to edge gateways rather than performing them on the sensing device directly.

DOI: 10.46121/pspc.54.2.66

Future work includes formal verification of the hybrid handshake protocol using automated protocol-analysis tools; extending the framework with post-quantum secure firmware update and attestation using SLH-DSA; integrating confidential-computing enclaves at the cloud tier to protect the PQ-KMS root of trust; and staging a live pilot deployment within a partner hospital and university campus to validate the framework's performance and operational characteristics under real clinical and facilities workloads. As NIST's fourth-round selection of HQC and the forthcoming FIPS 206 (FN-DSA/Falcon) standard mature, the framework's crypto-agile key-management design is intended to allow these additional algorithms to be incorporated without re-architecting the underlying digital twin security model, ensuring that healthcare and smart campus operators can adapt as the post-quantum cryptographic landscape continues to evolve.

REFERENCES

1. **Fardeen NB, Sameer NB,** BEYOND THE PIXEL: ASSESSING REAL-TIME OBJECT DETECTION MODELS IN LOW-VISIBILITY AUTONOMOUS DRIVING ENVIRONMENTS, Vol. 53 No. 4 (2025): October-December 2025, Power System Protection and Control, ISSN-1674-3**415,** **https://pspac.info/index.php/dlbh/article/view/371**
2. **Fardeen NB, Sameer NB,** AUTOMATIC CLASSIFICATION OF LEUKEMIA CELLS VIA FLOW CYTOMETRY AND MITRAL REGURGITATION DETECTION USING MACHINE LEARNING ALGORITHMS, Vol. 49 No. 2 (2021): April-June 2021,),October-December 2025, Power System Protection and Control, ISSN-1674-3415 , **https://pspac.info/index.php/dlbh/article/view/361**
3. **Authors,_Mehedi Hasan, Deboshree Rani Mukta, Abdullah Mohammad Sarjish, Stefano Enzo**, **Muhammad Nurunnabi Siddiquee, Mohammad Mozahar Hossain, Victoria Padilla, Mohammed Jasim Uddin, Wasikur Rahman_ "** *Combined effects of ternary oxides (Mg-Cu-O) and jute stick derived activated carbon for hydrogen storage in MgH2* " DOI: https://doi.org/10.1007/s11356-025-36394-4 , **Apr-2025**
4. **Authors: _Md. Wasikur Rahman Md. Sabbir Ahmed, Selim Reja, Sk Md Ali Zaker Shawon, Farhana Nazmin, Abdullah Mohammad Sarjish, Yousuf Ali, M. Azizur R. Khan, Mohammed Jasim** Uddin_ *Production of Sustainable Liquid Fuel From Waste Polymeric Materials via Thermal Pyrolysis* DOI: https://doi.org/10.1155/er/1449087 ,
5. **Sarjish, Abdullah Mohammad, et al.** " *Nano Polymers for Clean Energy Solutions* ." Nanostructured Polymers: Synthesis and Performance. Singapore: Springer Nature Singapore, 2026. 461-495. https://doi.org/10.1007/978-981-95-3581-1_18 , **Feb-2026**
6. **Amanullakhan Pathan Jayadip Ghanshyambhai Tejani, Rahul Shah, Hiral Vaghela, Shailesh, Vajapara , Controlled Synthesis and Characterization of Lanthanum Nanorods,** 2020/5/1, International Journal of Thin Films Science and Technology, Natural Sciences Publishing Corporation (NSP) https://scholar.google.com/citations?view_op=view_citation&hl=en&user=efbNMkwAAAAJ&citation_for_view=efbNMkwAAAAJ:M3NEmzRMIkIC
7. **Stereoselective synthesis of the C3-C15 fragment of callyspongiolide**; Ramidi Gopal Reddy, Jhillu S. Yadav and Debendra K. Mohapatra. (Tetrahedron Letters. 2018, 59, 3579-3582). https://doi.org/10.1016/j.tetlet.2018.08.041, https://www.sciencedirect.com/science/article/pii/S0040403918310426
8. **Asymmetric Total Synthesis of Two Possible Diastereomers** of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063).

DOI: 10.46121/pspc.54.2.66

https://doi.org/10.1021/acs.joc.6b02611,
https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611

9. **Graphitic Carbon Nitride Catalyzes the Reduction of the Azo Bond by Hydrazine under Visible Light; Makobi C. Okolie**, Glory G. Ollordaa, Gopal R. Ramidi, Xin Yan, Yufeng Quan, Qingsheng Wang and Yingchun Li. (Nanomaterials 2024, 14(17), 1402), https://doi.org/10.3390/nano14171402, https://www.mdpi.com/2079-4991/14/17/1402
10. **Lakshmi Rahul Reddy Mareddy, 2026, IJCNIS,** https://www.ijcnis.org/index.php/ijcnis/article/view/8735
11. **Lakshmi Rahul Reddy Mareddy, 2026, ICAA, https://eudoxuspress.com/index.php/pub/article/view/4780/3569**
12. **Lakshmi Rahul Reddy Mareddy, 2026, IJCMI , https://www.ijcmi.in/index.php/ijcmi/article/view/70**
13. **Rangavinay Teja Royal Jagga** , 2026, https://ieeexplore.ieee.org/document/11448495
14. **Rangavinay Teja Royal Jagga , 2026, https://ieeexplore.ieee.org/document/11448514**
15. **Rangavinay Teja Royal Jagga , 2026, https://ieeexplore.ieee.org/document/11448337**
16. **Rangavinay Teja Royal Jagga , 2026, https://ieeexplore.ieee.org/document/11497560**
17. **Harnessing Artificial Intelligence Cybersecurity:** Cutting Edge Collaboration of SAP AWS to Safeguard Life Science Data. 2026 https://doi.org/10.1109/aiei69164.2026.11497833
18. **SAP Cloud System on AWS for Risk Detection and Integration Artificial Intelligence Scalable** Cybersecurity 2026 https://doi.org/10.1109/aiei69164.2026.11497435
19. **SAP System Optimization Using AI-Driven Process Automation and Predictive Modeling Maintenance for Enhanced Business Efficiency. 2025** https://doi.org/10.1109/computingcon64838.2025.11376613
20. https://ieeexplore.ieee.org/abstract/document/11376613
21. https://ieeexplore.ieee.org/abstract/document/11497833
22. https://ieeexplore.ieee.org/abstract/document/11497435
23. **Kaleshwar Aryasomayajula**, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/294
24. **Kaleshwar Aryasomayajula, MODERN** DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/364
25. **Kaleshwar Aryasomayajula**, Micro services: "A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments., Vol. 51 No. 2 (2023): April-June 2023 , Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/374
26. **Alam, M. Z., Rahman, R., Sozib, H. M., Ahmed**, H., Hossain, A., Sabeena, A. A., Tasnim, A. F., Ahmed, F., Sarkar, M. I., & Erdei, T. I. (2026). Enhancing Thyroid Disease Diagnosis with

DOI: 10.46121/pspc.54.2.66

Machine Learning and Counterfactual Explainable AI. IEEE Access, 1–1. https://doi.org/10.1109/access.2026.3663497

27. **Yeasmin, S., Semi, M. M. A., Rony, M. K. K., Das, S., Sabeena, A. A., Rahman, R., Biswas, B., Ahmed, F., & Hossain, A. (2025).** Artificial Intelligence for Mental Health Monitoring: A Solution for Digital Behavioral Health Care and Education—An Umbrella Review. Health Science Reports, 9(1), e71703. https://doi.org/10.1002/hsr2.71703
28. **Hasan, S., Rahman, K. A., Ahmed, F., & Hossain, A. (2026). An integrated AI-driven framework for** maternal resource intelligence shortages across U.S. hospitals. Integrative Biomedical Research, 10(1), 1–8. https://doi.org/10.25163/biomedical.10110694
29. **Riipa, M. B., Ahmed, F., Rony, M. K. K., Hossain, A., Islam, A., Utsho, M. R., Kamal, M. B., Sharmin, S., &** Tasnim, A. F. (2026). The role of artificial intelligence in predicting cardiovascular outcomes: a systematic review and meta-analysis. Biostatistics & Epidemiology, 10(1). https://doi.org/10.1080/24709360.2026.2670804
30. **Semi, M. M. A., Das, S., Utsho, M. R., Hossain,** A., Kamal, M. B., Sizan, A. A., Tasnim, A. F., Yeasmin, S., & Parvin, M. R. (2026). Artificial Intelligence in Public Health Education: A Scoping Review of workforce competency development. Health Science Reports, 9(3). https://doi.org/10.1002/hsr2.72066
31. **Ahmed, F., Hasan, S., Hossain, A., & Rahman, K. A. (2026). Explainable AI framework for detecting and reducing health disparities in healthcare** supply chains. Journal of AI, ML and DL, 2(1), 1–8. https://doi.org/10.25163/ai.2110685
32. **Godavari Modalavalasa," SCALABLE CLOUD SOLUTIONS THROUGH ARTIFICIAL INTELLIGENCE GOVERNANCE**: APPLICATIONS IN HEALTHCARE AND FINANCIAL SYSTEMS", Vol. 54 No. 1 (2026): January-March 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/219, , DOI: https://doi.org/10.46121/pspc.54.1.25
33. **Godavari Modalavala**sa, AUTONOMOUS DATA ENGINEERING PIPELINES: A POLICY-DRIVEN ARCHITECTURE FOR SECURE AND SCALABLE CLOUD-NATIVE ANALYTICS, Vol. 53 No. 4 (2025): October-December 2025, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/181, DOI: https://doi.org/10.46121/pspc.53.4.25
34. **Godavari Modalavalasa, LARGE LANGUAGE MODELS FOR INTELLIGENT DATA ENGINEERING: AUTOMATING** SCHEMA DESIGN, LINEAGE, AND QUALITY CONTROL, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/183 , DOI: https://doi.org/10.46121/pspc.50.2.4
35. **Godavari Modalavalasa**, FEDERATED LEARNING FOR ENTERPRISE CLOUD DATA ENGINEERING: ARCHITECTURE, SECURITY, AND GOVERNANCE CHALLENGES, Vol. 51 No. 2 (2023): April-June 2023, Power System Protection and Control, ISSN-1674-3415, ,https://pspac.info/index.php/dlbh/article/view/184, DOI: https://doi.org/10.46121/pspc.51.2.5
36. **Godavari Modalavalasa, AI**-DRIVEN DATA GOVERNANCE: INTELLIGENT METADATA, LINEAGE, AND COMPLIANCE AUTOMATION IN CLOUD DATA PLATFORMS, Archives / Vol. 52 No. 1

(2024): January-March 2024 /, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/182, DOI: https://doi.org/10.46121/pspc.52.1.3

37. **Hima Bindu Lekkala, VishnuVardhan Bandar**i , AUTONOMOUS WORKFLOW OPTIMIZATION USING MULTI AGENT AI SYSTEMS AI AGENTS MANAGE STATIONS, WIP, AND TASK HANDOFFS, Vol. 54 No. 2 (202): April-June 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/306 , DOI: https://doi.org/10.46121/pspc.54.2.08
38. **Lekkala, H. B., & Bandari, V. V. (2026). Autonomous workflow optimization using multi-agent AI systems:** AI agents manage stations, WIP, and task handoffs. *Power System Protection and Control*, 54(2), 95–101. http://pspac.info/index.php/dlbh/article/view/306
39. **Bandari, V. V., & Lekkala, H. B. (2024). Physics-informed reinforcement learning for real-time control of complex manufacturing processes**. *Power System Protection and Control*, 52(2), 164–170. https://pspac.info/index.php/dlbh/article/view/343
40. **Lekkala, H. B., & Bandari, V. V. (2025). AI-based energy-aware scheduling and process optimization** in engineer-to-order smart manufacturing systems. *Power System Protection and Control*, 53(1), 30–36. http://pspac.info/index.php/dlbh/article/view/341
41. **Bandari, V. V., & Lekkala, H. B. (2026). AI-driven digital** thread framework for end-to-end lifecycle optimization in ETO manufacturing systems. *Power System Protection and Control*, 54(2), 318–325. http://pspac.info/index.php/dlbh/article/view/340
42. **Lekkala, H. B., & Bandari, V. V. (2026). Deep learning** for weld defect detection using CNN or vision transformers fusion detection. *Power System Protection and Control*, 54(2), 151–158. https://pspac.info/index.php/dlbh/article/view/314
43. **Bandari, V. V., & Lekkala, H. B. (2024). AI-based operator** behavior monitoring and cost optimization using digital traceability in manufacturing systems. *Power System Protection and Control*, 52(1), 38–45. https://pspac.info/index.php/dlbh/article/view/342
44. **Shreyas Jayaprakash,** MACHINE LEARNING-DRIVEN VERIFICATION: USING ML TO OPTIMIZE COVERAGE ANALYSIS, PREDICT SIMULATION RUNTIMES, AND AUTO-GENERATE TESTBENCHES, Vol. 54 No. 2 (2026): April-June 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/391
45. **Mohammed Shafi Kundiladi,** EVENT-DRIVEN IMAGE AND VEHICLE STATUS MANAGEMENT FOR LOW-POWER IOT DIGITAL LICENSE PLATES, Vol. 53 No. 3 (2025): July-September 2025, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/175 DOI: https://doi.org/10.46121/pspc.53.3.17
46. : Controlled Synthesis and Characterization of Lanthanum Nanorods, **Author(s), Jayadeep Tejani, Rahul Shah, Hiral Vaghela, Shailesh Vajapara, Amanullakhan Pathan**, doi:10.18576/ijtfst/090205, URL: https://www.naturalspublishing.com/Article.asp?ArtcID=20705, May-2020
47. **Kunal Kapoor**, **Apr 2026,** Relation Extraction and NER through Fine Tuned BERT model with transfer Learning, https://ieeexplore.ieee.org/document/11472214/metrics#metrics

DOI: 10.46121/pspc.54.2.66

48. **Harender Bisht**, Human–AI Collaboration in Insurance Fraud Detection: Ethical Cloud-Native Architectures for Fair and Transparent Decision Support, **Volume 2026, Issue 1,** https://doi.org/10.52710/cfs.873
49. **Harender Bisht,** Ethical AI-Augmented Compliance Systems: Architectural Innovations for Cloud-Native Financial and Insurance Data Integrity, **Vol. 11 No. 1s (2026) ,** https://doi.org/10.52783/jisem.v11i1s.14056
50. **Harender Bisht**, Governance-By-Design For AI-Based Insurance Fraud Detection: Auditability, Accountability, And Regulatory Traceability, **Archives / Vol. 9 No. 1 (2026**) ,https://doi.org/10.63278/jicrcr.vi.3620
51. **Harender Bisht,** Operationalizing Explainable AI in Insurance Fraud Detection: From Model Transparency to Decision Justification, **Vol. 35 No. 1 (2026): JOCAAA**, https://eudoxuspress.com/index.php/pub/article/view/4791,
52. **A. MahendraVardhan and S. Sridhar,** "Determining False Positive Analysis of Software Vulnerabilities with Predefined Scan Rules using Random Forest Classifier and Decision Tree Technique," **2022 4th International Conference on Advances in Computing**, Communication Control and Networking (ICAC3N), Greater Noida, India, 2022, pp. 622-625, **DOI: https://doi.org/10.1109/ICAC3N56670.2022.10074458**
53. **Amilineni, M.V. 2025**. Audit-Safe Agentic RAG Framework for Regulated Enterprise Systems: A Trustworthy AI Architecture for SAP, Finance, Healthcare, and Government Workflows. INTERNATIONAL JOURNAL OF ADVANCES IN SIGNAL AND IMAGE SCIENCES. (Jan. 2025), 34–46. **DOI: https://doi.org/10.29284/7fjzk883**
54. **Islam, Imtiaz (Annotator), Goard, Michael (Annotator), Daves, Daniel (Annotator),** Credit Card Fraud Detection: A Comparative Study of Anomaly Detection and Supervised Learning under Extreme Class Imbalance, Published **May 5, 2026** | Version v1, **DOI: 10.5281/zenodo.20032207**
55. **Deepa Nagalavi; Shankar Balla Microsoft, USA; Pushpalatha. M; Nashwan Adnan Othman; Malik Bader Alazzam; A. Balakumar, Enhancing Real-Time Language Processing via Advanced PET Signal Analysis and Deep Learning,** ***06-07 June 2025***, DOI: 10.1109/ICICKE65317.2025.11136561 , https://ieeexplore.ieee.org/document/11136561
56. **Venkata Sai Aditya Kondru,FMR, Carmel, USA ,; Shankar Balla; Dhavalkumar Thakar; Dheeraj Velaga; Sri Lekha Bandla,** Intelligent Human–Computer Interaction for Navigation Control Through Vision-Based Hand Gesture Recognition, Conferences >2026 IEEE 15th International , **Date of Conference:** ***07-09 April 2026***,**Date Added to IEEE** ***Xplore*****:** 08 May 2026 **DOI:** 10.1109/CSNT69054.2026.11502317 , https://ieeexplore.ieee.org/abstract/document/11502317
57. **Sudipkumar Ghanvat , Aishwarya Badlani, Shreya Sandeep Joshi,** Marketing Effectiveness in the Post-Cookie Era: A Comparative Analysis of First Party and Zero-Party Data Strategies on Campaign Performance and Customer Equity, **Vol. 31 No. 4 (2023):** JOCAAA , https://www.eudoxuspress.com/index.php/pub/article/view/3940
58. **Navin Chhibber,; Deepak Singh; Anokh Kishore, Capital One, USA; Nikita Chawla; K. Anguraj,** Multi-Granular Attention-Driven Reinforcement Learning Framework for Web Intelligent Enhancement Systems, , **11 June 2026**, https://ieeexplore.ieee.org/document/11483386

59. **Kapoor K, et al. AI-Driven COVID-19 Care**: Analytics, Scheduling & Geo-Access. International Journal of Drug Delivery Technology (IJDDT). **2026;16(65S):Article** 136. Available at: https://impactfactor.org/PDF/IJDDT/16/IJDDT,Vol16,Issue65s,Article136.pdf
60. **Kilama, D. (2026). Mechanisms of Chinese** Aid to Africa: A Comparative Analysis with Western Philanthropic and Development Models. Interdisciplinary Social Research, 2(1).DOI: https://doi.org/10.64220/isr.v2i1.011, https://scholarlysummit.com/journals/isr/articles/mechanisms-of-chinese-aid-to-africa-a-comparative-analysis
61. **Kilama, D. (2026). Beyond Altruism**: Analyzing China's Strategic Motivations Behind Development Aid to Africa. *Interdisciplinary Social Research, Volume 2 (2026)*(Issue 1), DOI: . https://doi.org/10.64220/isr.v2i1.012, Beyond Altruism: Analyzing China's Strategic Motivations Behind Development Aid to Africa
62. Anumolu DPK, Veena B, Afreen SS, Bandla J, Lakshmi KC, Pulusu VS. In vitro models for studying drug metabolites and metabolizing enzymes: A comprehensive review. Int J Pharm Investig. 2026;16(3):885–91. Available from: http://dx.doi.org/10.5530/ijpi.20260116
63. Asgar Z, Kumar G, Khandelwal P, Pratap B, Gurjar V, Baluni A, et al. Utility of contrast-enhanced computed tomography abdomen in Intestinal Obstruction. Int J Drug Deliv Technol. 2026;16(32s). Available from: http://dx.doi.org/10.25258/ijddt.16.32s.4
64. Anumolu DP, Ramatulasi J, Ashok G, Afreen SS, Mathew C, Shakar PV. Synthesis and characterization of MIP ghost approach for selective extraction of empagliflozin and quantification by liquid chromatography. J Chromatogr Sci. 2025;63(2). Available from: http://dx.doi.org/10.1093/chromsci/bmaf008.
65. Anumolu DPK, Naraparaju S, Addada SB, Pagidipally V, Pulusu VS, Afreen SS. Synthesis of silver nanoparticles using Ecbolium ligustrinum leaves: Characterization through advanced analytical techniques. Nanotechnol Percept . 2024;710–8. Available from: http://dx.doi.org/10.62441/nano-ntp.vi.3618
66. Naraparaju S, Mukti A, Anumolu DP, Chaganti S. Development and validation of a spectrofluorimetric method for the quantification of capecitabine in bulk and tablets. Turk J Pharm Sci [Internet]. 2023;20(4):234–9. Available from: http://dx.doi.org/10.4274/tjps.galenos.2022.46364
67. Naraparaju S, Yamjala P, Chaganti S, Anumolu DP. Spectrophotometric quantification of atomoxetine hydrochloride based on nucleophilic substitution reaction with 1,2-naphthoquinone-4-sulfonic acid sodium salt (NQS). Turk J Pharm Sci [Internet]. 2024;20(6):405–11. Available from: http://dx.doi.org/10.4274/tjps.galenos.2022.09147.
68. Anumolu PD, Shakar PV, Tulasi JR, Soujanya C, Afreen SS, Ashok G. Concurrent discriminative emission intensity quantification of Fexofenadine hydrochloride and Montelukast Sodium. Orient J Chem [Internet]. 2022;38(6):1364–8. Available from: http://dx.doi.org/10.13005/ojc/380605
69. Anumolu PD, Gurrala S, Menkana S, Mathew C, Srimantula L. Application of fluorotag in quantification of metformin in tablet dosage form and biological samples. The Thai Journal of Pharmaceutical Sciences [Internet]. 2022;46(1):93–8. Available from: http://dx.doi.org/10.56808/3027-7922.2550

70. Gurrala S, Raj S, Cvs S, Anumolu PD. Quality-by-design approach for chromatographic analysis of metformin, empagliflozin and linagliptin. J Chromatogr Sci [Internet]. 2022;60(1):68–80. Available from: http://dx.doi.org/10.1093/chromsci/bmab030

71. Gurrala S, Raj S, Cvs S, Anumolu DP, Naraparaju S, Nizampet H. Response surface methodology in spectrophotometric estimation of saxagliptin, derivatization with MBTH and ninhydrin. Turk J Pharm Sci [Internet]. 2022;19(1):9–18. Available from: http://dx.doi.org/10.4274/tjps.galenos.2021.93195

72. National Institute of Standards and Technology, "Module-Lattice-Based Key-Encapsulation Mechanism Standard," FIPS 203, Aug. 2024.

73. National Institute of Standards and Technology, "Module-Lattice-Based Digital Signature Standard," FIPS 204, Aug. 2024.

74. National Institute of Standards and Technology, "Stateless Hash-Based Digital Signature Standard," FIPS 205, Aug. 2024.

75. National Institute of Standards and Technology, "NIST Selects HQC as Fifth Algorithm for Post-Quantum Encryption," NIST News, Mar. 2025.

76. D. Moody et al., "Status Report on the Fourth Round of the NIST Post-Quantum Cryptography Standardization Process," NIST Internal Report NIST IR 8545, 2025.

77. National Institute of Standards and Technology, "Recommendations for Key-Encapsulation Mechanisms," Special Publication 800-227 (Draft), 2025.

78. P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," SIAM J. Comput., vol. 26, no. 5, pp. 1484–1509, 1997.

79. M. Mosca, "Cybersecurity in an era with quantum computers: Will we be ready?," IEEE Security & Privacy, vol. 16, no. 5, pp. 38–41, 2018.

80. O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," J. ACM, vol. 56, no. 6, pp. 1–40, 2009.

81. J. Bos et al., "CRYSTALS-Kyber: A CCA-secure module-lattice-based KEM," in Proc. IEEE Eur. Symp. Security and Privacy (EuroS&P), 2018, pp. 353–367.

82. L. Ducas et al., "CRYSTALS-Dilithium: A lattice-based digital signature scheme," IACR Trans. Cryptogr. Hardw. Embed. Syst., vol. 2018, no. 1, pp. 238–268, 2018.

83. D. J. Bernstein et al., "SPHINCS+: Practical stateless hash-based signatures," in Proc. EUROCRYPT, 2015, pp. 368–397.

84. M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in Transdisciplinary Perspectives on Complex Systems, Springer, 2017, pp. 85–113.

85. F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," IEEE Trans. Ind. Informat., vol. 15, no. 4, pp. 2405–2415, 2019.

86. E. Glaessgen and D. Stargel, "The digital twin paradigm for future NASA and U.S. Air Force vehicles," in Proc. 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conf., 2012.

87. M. A. Khan, M. Y. Ismail, and R. Iqbal, "A survey of digital twin security: Threats, challenges, and open research directions," IEEE Access, vol. 11, pp. 34120–34146, 2023.

DOI: 10.46121/pspc.54.2.66

88. S. Erol, A. Wallmuller, and W. Narodoslawsky, "Digital twin applications in healthcare: A systematic review," IEEE J. Biomed. Health Informat., vol. 27, no. 6, pp. 2678–2690, 2023.
89. R. Minerva, G. M. Lee, and N. Crespi, "Digital twin in the IoT context: A survey on technical features, scenarios, and architectural models," Proc. IEEE, vol. 108, no. 10, pp. 1785–1824, 2020.

DOI: 10.46121/pspc.54.2.66