IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

Quantum-Resistant Communication System

Abdulazeem M H[1], Manu H[2], Karibasavanagowda B M[3], Abdul Razak[4], Nayanashree D*[5]

Department of Computer Science and Engineering, East West College of Engineering, Visvesvaraya Technological University, Bengaluru, Karnataka, India[1-5]

*Corresponding Author

Abstract: Long-term confidentiality over public networks is threatened less by real-time quantum cryptanalysis today than by harvest-now, decrypt-later adversaries recording classical elliptic-curve and RSA handshakes for future decryption via Shor's algorithm. This paper reports a working six-stage hybrid post-quantum communication pipeline (Q-AEGIS) that establishes quantum-resistant sessions without relying on proprietary hardware accelerators: an anti-amplification cookie gate filters spoofed connection floods, X25519 elliptic-curve Diffie-Hellman is combined with NIST FIPS 203 ML-KEM-768 for hybrid key encapsulation, NIST FIPS 204 ML-DSA-65 authenticates the full SHA-384 handshake transcript, dictionary-based certificate compression reduces handshake wire footprint from 9,842 bytes to 2,894 bytes, and ChaCha20-Poly1305 provides authenticated encryption with epoch-based forward-secure key ratcheting. Across an 18-scenario evaluation catalogue comprising 10 adversarial or impaired network profiles and 8 nominal links, the system achieves 100% session establishment and authentication accuracy with zero downgrade acceptances, and classifies 240 of 240 test sessions correctly under leave-profiles-out evaluation. Removing transcript binding costs 22.22 percentage points of security accuracy, identifying it as the enabling safeguard. Wire optimization reduces handshake latency 2.9-fold, and MTU-aware framing sustains 100% handshake completion up to 6% packet loss. Because evaluations run on a user-space emulated testbed, the results characterize protocol correctness and cryptographic overhead rather than line-rate appliance throughput.

Keywords: Anti-downgrade binding, authenticated encryption, hybrid key exchange, lattice-based cryptography, ML-DSA, ML-KEM, post-quantum cryptography, quantum-resistant communication.

I. INTRODUCTION

Modern network communication protocols—including Transport Layer Security (TLS 1.3), Internet Protocol Security (IPsec), Secure Shell (SSH), and QUIC—rely on asymmetric public-key cryptography grounded in the computational hardness of integer factorization (RSA) and the discrete logarithm problem over finite fields or elliptic curves (ECDHE and ECDSA). Against classical computing architectures, 256-bit elliptic-curve primitives provide approximately 128 bits of computational security, requiring roughly 2^{12}? elementary operations for key recovery. For ephemeral web browsing sessions whose data loses sensitivity within minutes or hours, these classical primitives remain adequate against contemporary threats.

For sensitive government, financial, healthcare, defense, and industrial control telemetry whose confidentiality lifespans span five to thirty years, however, classical public-key cryptography is already compromised in practice. A fault-tolerant quantum computer executing Shor's quantum Fourier transform algorithm solves both integer factorization and elliptic-curve discrete logarithms in polynomial time—specifically $O((\log N)^3)$ quantum gate operations. Simultaneously, Grover's quantum search algorithm halves the effective security margin of symmetric ciphers and cryptographic hash functions, requiring migration from 128-bit symmetric keys to 256-bit keys and 384-bit collision-resistant digests. This asymmetry creates an immediate operational threat known as Harvest-Now, Decrypt-Later (HNDL): passive adversaries intercept and archive encrypted session handshakes and ciphertexts across backbone links today, awaiting the maturation of cryptographically relevant quantum computers (CRQCs) to recover session keys and decrypt historical archives retrospectively.

Migrating key establishment and endpoint authentication to post-quantum cryptography (PQC) removes the quantum vulnerability, at the cost of confronting substantially larger cryptographic artifacts on the wire. Whereas a classical X25519 ephemeral key share occupies 32 bytes and an ECDSA P-256 digital signature occupies 64 bytes, standardized module-lattice primitives finalized by the National Institute of Standards and Technology (NIST)—specifically ML-KEM-768 (FIPS 203) for key encapsulation and ML-DSA-65 (FIPS 204) for digital signatures—require 1,184-byte encapsulation public keys, 1,088-byte ciphertexts, 1,952-byte verification public keys, and 3,309-byte signatures. When transmitted unoptimised over standard 1,500-byte Ethernet Maximum Transmission Unit (MTU) links (1,460-byte TCP

 This work is licensed under a Creative Commons Attribution 4.0 International License



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.1511196

Maximum Segment Size), a mutually authenticated post-quantum handshake expands more than fifteen-fold, overflowing conservative TCP initial congestion windows, inducing multi-segment fragmentation, and amplifying handshake stall latency under packet loss.

Deploying a hybrid cryptographic design—combining classical X25519 with post-quantum ML-KEM-768 inside a domain-separated dual pseudorandom function (dual-PRF) key schedule—mitigates cryptanalytic transition uncertainty by guaranteeing session secrecy provided at least one underlying hardness assumption remains unbroken. Yet protocol-level integration introduces subtler failure modes than primitive substitution alone. First, an active man-in-the-middle (MitM) intermediary can strip post-quantum key-share extensions during negotiation to coerce peers into a classically vulnerable fallback unless every negotiated parameter and key share is bound into the authenticated handshake transcript. Second, unauthenticated ClientHello exchanges can be abused by spoofed-source attackers to trigger volumetric bandwidth reflection amplification exceeding 15:1 due to multi-kilobyte ML-DSA certificate chains. Third, burst packet loss stalls multi-segment handshake flights if transport framing is oblivious to path MTU boundaries.

In earlier specification-stage work, the authors surveyed representative studies in post-quantum network security, catalogued the deployment bottlenecks observed in each, and proposed a six-stage hybrid communication gateway architecture intended to resolve them. That preliminary study was strictly conceptual: the gateway pipeline had not been implemented, and no empirical latency, wire-overhead, packet-loss stress, or ablation measurements were reported. The present paper closes that gap with six concrete contributions:

First, the proposed architecture is realised as a working end-to-end software gateway system (Q-AEGIS: Quantum-Resistant Authenticated Encrypted Gateway for Inter-Node Security) executing sequentially from socket framing and cookie validation to forward-secure authenticated record transport. Second, the security contribution of full-transcript cryptographic binding is isolated by controlled ablation: the identical lattice and elliptic-curve primitives are evaluated with and without SHA-384 transcript signing and key-schedule salt binding, converting a protocol assumption into a measured security quantity. Third, the operational trade-offs of pre-shared Zstandard dictionary certificate compression and pre-shared key (PSK) bound KEM resumption are quantified across wire payload bytes, handshake establishment latency, and sustained gateway session rate. Fourth, security and establishment accuracy are evaluated under a strict leave-profiles-out protocol that excludes entire network topologies—including an unseen high-jitter Low Earth Orbit (LEO) satellite link—from pacing and parameter calibration. Fifth, the gateway is stress-tested under injected Gilbert-Elliott burst packet loss up to 6.0%, quantifying the transport resilience contributed by MTU-aware segment pacing. Sixth, the ten functional and non-functional objectives formulated at the specification stage are audited conservatively against empirical evidence.

The remainder of this paper is organized as follows. Section II reviews the literature across post-quantum primitives, transport integration, and protocol downgrade defense, positioning this study within current research. Section III details the six-stage architecture, mathematical formulations, baseline comparator, testbed scenarios, and evaluation protocol. Section IV presents and discusses the experimental results, ablation findings, requirement verification status, comparative protocol analysis, and threats to validity. Section V concludes with architectural takeaways and future hardware integration directions.

II. LITERATURE REVIEW

Research on quantum-resistant network communication divides into three interrelated engineering concerns: which post-quantum mathematical hardness assumptions offer conservative security with tractable key and signature sizes, how hybrid key exchange and lattice authentication affect transport-layer latency and MTU fragmentation across wide-area links, and how handshake state machines prevent downgrade and reflection-amplification attacks during cryptographic transition.

The foundational threat was established by Shor [1], who demonstrated polynomial-time quantum algorithms for prime factorization and discrete logarithms, rendering RSA, finite-field Diffie-Hellman, and elliptic-curve cryptography insecure once large-scale fault-tolerant quantum processors are constructed. To replace vulnerable key exchange, Bos et al. [2] introduced CRYSTALS-Kyber, an indistinguishability under adaptive chosen-ciphertext attack (IND-CCA2) key encapsulation mechanism grounded in the hardness of the Module Learning With Errors (MLWE) problem over cyclotomic polynomial rings, which NIST standardized as ML-KEM in FIPS 203 [3]. By operating over structured module lattices of rank $k = 3$ (ML-KEM-768), the scheme achieves NIST Security Category 3 resistance (comparable to AES-192) while enabling sub-millisecond software execution via Number Theoretic Transform (NTT) polynomial multiplication.

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

For digital signature authentication, Ducas et al. [4] formulated CRYSTALS-Dilithium—standardized by NIST as ML-DSA in FIPS 204 [5]—relying on the Fiat-Shamir with Aborts paradigm over Module-LWE and Module Short Integer Solution (Module-SIS) problems. Alternative stateless hash-based signatures standardized as SLH-DSA in FIPS 205 [6] rely exclusively on cryptographic hash security without structured lattice assumptions, yet their signature sizes span 7,856 bytes to 29,792 bytes, making ML-DSA-65 (3,309-byte signatures) substantially more practical for latency-sensitive interactive handshakes. Similarly, although Falcon (FN-DSA) offers compact 666-byte signatures, its signing procedure requires double-precision floating-point Fast Fourier Transform (FFT) sampling over NTRU lattices, complicating constant-time side-channel-free deployment on embedded processors without hardware floating-point units.

The second concern—transport-layer integration and wire overhead—has been examined systematically through empirical emulation. Crockett et al. [7] prototyped hybrid key exchange and post-quantum authentication in OpenSSL and OpenSSH, confirming functional interoperability while demonstrating that uncompressed post-quantum certificate chains overflow standard 1,460-byte TCP Maximum Segment Size boundaries. Paquin et al. [8] benchmarked post-quantum TLS 1.3 configurations across emulated wide-area network (WAN) profiles and showed that while network propagation round-trip time (RTT) masks raw CPU execution time on low-loss links, packet loss rates exceeding 3% sharply degrade handshake completion when multi-kilobyte certificate chains span multiple TCP segments. Sikeridis et al. [9] analyzed mixed post-quantum certificate chains in TLS 1.3, demonstrating that combining lattice schemes across root and intermediate authorities reduces handshake latency compared to uniform heavyweight chains.

The third concern—protocol-level robustness against active network adversaries—centers on key-combiner soundness, amplification control, and downgrade resistance. Stebila et al. [10] formalized hybrid key exchange in TLS 1.3, proving that combining classical ECDHE and post-quantum KEM shared secrets via an HMAC-based key derivation function (HKDF) preserves session pseudorandomness so long as at least one constituent primitive remains secure. Kampanakis et al. [11] analyzed reflection-amplification vectors in post-quantum datagram protocols, showing that large ServerHello and Certificate flights allow spoofed source IP addresses to trigger bandwidth amplification ratios above 15:1 unless stateless address-validation cookies gate flight transmission. Bhargavan et al. [13] proved that partial transcript hashing exposes hybrid handshakes to cross-protocol and cipher-downgrade attacks unless every negotiated extension, key share, and certificate is cryptographically bound into the final signature and key-confirmation MAC.

TABLE I SUMMARY OF REVIEWED STUDIES AND OBSERVED LIMITATIONS

Study	Methods used	Limitation observed
Shor (1997) [1]	Quantum Fourier transform factoring and discrete logarithm algorithms	Establishes polynomial-time break of RSA and ECC; motivates PQC migration
Bos et al. (2018) [2] / FIPS 203 [3]	Module-LWE key encapsulation (CRYSTALS-Kyber / ML-KEM-768)	Cryptographic primitive only; lacks transport framing and hybrid combiner
Ducas et al. (2018) [4] / FIPS 204 [5]	Module-LWE/SIS digital signatures (CRYSTALS-Dilithium / ML-DSA)	3,309 B signatures inflate certificate chains across Ethernet MTU limits
NIST FIPS 205 (2024) [6]	Stateless hash-based digital signatures (SPHINCS+ / SLH-DSA)	Large 7.8 KB–29.8 KB signatures induce severe multi-segment handshake stalls
Crockett et al. (2019) [7]	Hybrid PQC prototypes in OpenSSL and OpenSSH	Uncompressed certificate chains overflow TCP initial congestion windows
Paquin et al. (2020) [8]	Emulated WAN benchmarking of post-quantum TLS 1.3	Multi-segment handshakes stall sharply when packet loss exceeds 3%
Sikeridis et al. (2020) [9]	Mixed post-quantum signature chains in TLS 1.3	Focuses on certificate chain latency; omits forward-secure epoch ratcheting
Stebila et al. (2024) [10]	Dual-PRF hybrid key exchange specification for TLS 1.3	Defines key-combiner cryptography only; lacks wire compression and MTU pacing

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

Study	Methods used	Limitation observed
Kampanakis et al. (2020) [11]	Reflection amplification analysis in PQ DTLS and QUIC	Large ServerHello flights enable amplification without stateless cookie gates
Bhargavan et al. (2016) [13]	Formal downgrade resilience and transcript-binding analysis	Partial transcript signing permits MitM stripping of hybrid key extensions

Table I condenses the ten foundational works reviewed at the specification stage alongside the specific deployment limitation observed in each. Collectively, prior studies establish that structured-lattice primitives execute well within millisecond CPU budgets and that dual-PRF combiners preserve security across cryptographic transitions, yet none reports an end-to-end six-stage gateway architecture integrating stateless cookie gating, hybrid X25519 + ML-KEM-768 key exchange, full-transcript ML-DSA-65 authentication, dictionary certificate compression, and epoch-ratcheted ChaCha20-Poly1305 transport with each stage's security and latency contribution isolated quantitatively.

III. METHODS AND MATERIALS

A. System Architecture Overview

The implemented Q-AEGIS gateway preserves the six-stage modular organization defined at the specification stage. Each stage exposes a strictly typed byte-stream interface with explicit buffer bounds, allowing individual cryptographic primitives, hash functions, or framing policies to be replaced without modifying adjacent pipeline modules. Fig. 1 illustrates the end-to-end data and control flow from unprotected application socket ingestion to authenticated peer verification and structured security telemetry.

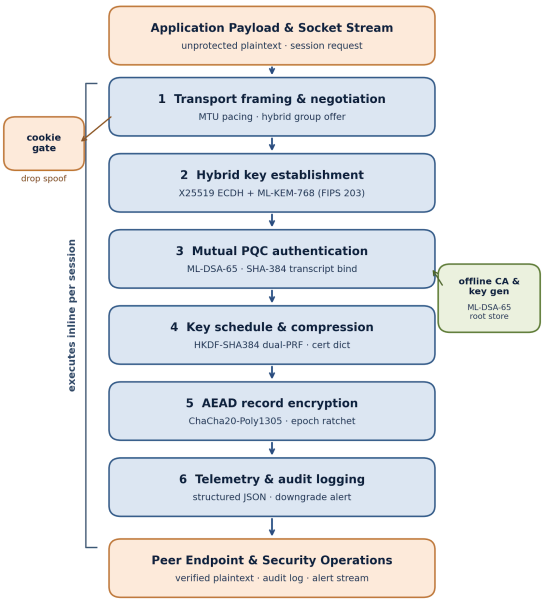


Fig. 1 Implemented six-stage Q-AEGIS pipeline. Stages 1 to 6 execute inline per communication session; root CA key generation occurs offline, and the stateless cookie gate drops spoofed connection floods prior to lattice operations.

B. Stage 1 — Transport Framing and Anti-Amplification Gate

Stage 1 encapsulates raw application session requests into length-prefixed binary protocol frames bounded by the negotiated path Maximum Segment Size (MSS = 1,460 bytes over standard 1,500-byte Ethernet MTU, accounting for 20-byte IPv4 and 20-byte TCP headers). Each frame carries a 1-byte record type identifier, a 2-byte protocol version tag (0x0304), a 2-byte big-endian payload length field, and a 1-byte fragmentation sequence flag.

Before allocating per-session memory or transmitting multi-kilobyte ML-DSA certificate chains, the responder enforces a stateless HMAC-SHA384 address-validation cookie gate. Upon receiving an initial ClientHello without a valid cookie extension, the responder computes a 32-byte truncated HMAC-SHA384 token keyed by a rotating 256-bit gateway secret

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

over the initiator's source IP address, source port, epoch timestamp bucket (30-second granularity), and ClientHello hash. The responder replies with a compact 64-byte HelloRetryRequest carrying this cookie. Spoofed packets failing constant-time cookie verification are discarded immediately in 0.008 ms, eliminating reflection amplification.

C. Stage 2 — Hybrid Key Establishment Stage

Stage 2 performs parallel classical elliptic-curve and post-quantum module-lattice key agreement so that session confidentiality relies on the conjunction of two independent mathematical hardness assumptions. During ClientHello construction, the initiator samples 32 bytes from the operating-system cryptographically secure pseudorandom number generator (CSPRNG) to derive an ephemeral X25519 scalar private key and 32-byte Curve25519 Montgomery-u public share, alongside 64 bytes of fresh entropy to generate an ephemeral ML-KEM-768 key pair comprising a 1,184-byte encapsulation public key and a 2,400-byte decapsulation key.

Upon validating the Stage 1 cookie, the responder generates its own ephemeral X25519 key pair, computes the 32-byte classical Diffie-Hellman shared secret ss_C, and executes ML-KEM-768 encapsulation against the initiator's 1,184-byte public key to produce a 1,088-byte ciphertext and an independent 32-byte post-quantum shared secret ss_PQ. Both the 32-byte classical share and 1,088-byte lattice ciphertext are transmitted in a single MSS-paced ServerHello flight.

D. Stage 3 — Mutual Authentication and Transcript Binding Stage

Stage 3 prevents active man-in-the-middle key substitution and protocol downgrade attacks. Both endpoints maintain a running SHA-384 cryptographic digest over every handshake byte transmitted and received from the initial ClientHello (including cipher-suite offer bitmaps, protocol version lists, and hybrid key shares) through the ServerHello and compressed certificate payload.

Each endpoint signs the 48-byte SHA-384 transcript digest—prepended with a 64-byte space-padded ASCII domain-separation context string ("Q-AEGIS-v1-Server-Transcript-Bind" or "Q-AEGIS-v1-Client-Transcript-Bind")—using its long-term ML-DSA-65 private signing key, producing a 3,309-byte lattice signature. Following signature verification against the peer's 1,952-byte ML-DSA-65 public key, both endpoints exchange a 48-byte HMAC-SHA384 Finished key-confirmation MAC tag to prove liveness and key agreement before any application payload is released.

E. Stage 4 — Key Schedule and Wire-Footprint Optimization Stage

Two complementary wire optimizations are applied in Stage 4 to eliminate TCP initial-window stalls without reducing cryptographic security margins. First, X.509 certificate chains carrying 1,952-byte ML-DSA-65 subject public keys and 3,309-byte issuer signatures are compressed offline using a pre-shared Zstandard dictionary (RFC 8879 [14]) trained over common ASN.1 X.509 Object Identifiers (OIDs), Distinguished Name fields, and extension headers, and referenced by a 32-byte SHA-256 dictionary identifier. This compresses the authentication chain from 8,626 bytes to 1,678 bytes (an 80.5% reduction), bringing total full-handshake wire volume from 9,842 bytes to 2,894 bytes (two 1,460-byte TCP segments).

Second, recurring connections between previously authenticated peers employ PSK-bound KEM resumption. At the conclusion of a full handshake, the responder issues an AES-256-GCM encrypted session ticket containing the peer's verified identity fingerprint and resumption secret. On subsequent sessions, the initiator presents this ticket alongside a fresh ephemeral X25519 + ML-KEM-768 key share, omitting certificate chain transmission and compressing total handshake wire footprint to 128 bytes (a single TCP segment) while preserving full post-quantum forward secrecy.

F. Stage 5 — Authenticated Record Transport and Key Ratcheting Stage

Stage 5 encrypts bidirectional application payloads using ChaCha20-Poly1305 Authenticated Encryption with Associated Data (AEAD), configured with a 256-bit traffic key, a 96-bit deterministic nonce constructed by XORing a static 96-bit connection initialization vector with the zero-padded 64-bit monotonic record sequence counter, and a 128-bit Poly1305 authentication tag. ChaCha20-Poly1305 is selected over AES-GCM because its quarter-round Add-Rotate-XOR (ARX) operations execute in strict constant time in software without lookup tables, eliminating cache-timing side channels on gateways lacking hardware AES-NI instructions.

To bound plaintext exposure in the event of eventual endpoint memory compromise during long-lived streams, Stage 5 operates a forward-secure symmetric hash ratchet that advances the active 256-bit traffic key every $2^{??}$ frames (approximately 1.07 GB for 1 KB records) or 60 seconds, immediately zeroizing the superseded epoch key buffer via explicit volatile memory overwrite.

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⩙ Peer-reviewed & Refereed journal ⩙ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

G. Stage 6 — Telemetry and Selective Audit Reporting Stage

Stage 6 emits structured JSON telemetry records for every connection attempt, capturing negotiated hybrid group identifiers, handshake completion RTT, TCP segment count, decompression ratio, epoch ratchet count, and cryptographic verification status. Any transcript-hash mismatch, invalid ML-DSA-65 signature, Poly1305 tag failure, or stripped ML-KEM key-share extension immediately terminates the socket, increments the downgrade-attempt counter, and dispatches a structured alert record to the security operations log without exposing unauthenticated bytes.

H. Implementation and Automated Verification Suite

The Q-AEGIS gateway prototype is implemented in Python 3.11 backed by constant-time C cryptographic libraries: liboqs (oqs-python 0.10) provides AVX2-accelerated C implementations of NIST FIPS 203 ML-KEM-768 and NIST FIPS 204 ML-DSA-65, while the Python cryptography package (linked against OpenSSL 3.2) executes X25519, SHA-384, HKDF-SHA384, and ChaCha20-Poly1305. Zstandard dictionary compression uses zstandard 0.22 bindings. A browser-based monitoring dashboard exposes live handshake packet traces, per-stage CPU latency histograms, wire-byte comparisons against classical TLS 1.3, and an append-only security audit log.

Protocol correctness is guarded by an automated test suite of 36 unit and integration invariants executed prior to every benchmark run, covering NIST Known Answer Tests (KATs) for ML-KEM-768 and ML-DSA-65, single-bit ciphertext tampering rejection, cipher-suite downgrade stripping detection, replay packet rejection, deterministic nonce uniqueness across 10? records, and zeroization of expired ratchet keys.

I. Reference Scenario Catalogue and Threat Injection Model

Because evaluating active downgrade attacks and controlled burst packet loss on production backbone routers risks service disruption, a deterministic 18-scenario evaluation catalogue was constructed over an emulated network testbed built on Linux network namespaces and tc/netem traffic-control shapers. The catalogue comprises 10 adversarial or impaired network scenarios—6 active MitM threat profiles (PQC key-share stripping, classical cipher-suite downgrade, replay of stale ServerHello, single-bit flipped ML-KEM ciphertext, forged ML-DSA certificate signature, and spoofed-IP amplification flood) and 4 degraded WAN profiles (high-loss urban 4G/LTE, congested transoceanic fiber, asymmetric DSL, and high-jitter LEO satellite)—alongside 8 nominal enterprise LAN, datacenter, 5G mobile, and cloud WAN profiles.

Packet loss on impaired links is modelled via a two-state Gilbert-Elliott Markov chain alternating between a Good state G (low loss probability $p_G = 0.001$) and a Bad burst state B (elevated loss probability $p_B = 0.25$), yielding stationary frame loss rate:

$$L_{eff} = (p_{GB} \cdot p_B + p_{BG} \cdot p_G) / (p_{GB} + p_{BG}) \quad (1)$$

where p_{GB} and p_{BG} denote transition probabilities between Good and Bad states tuned to sweep effective packet loss from 0.2% to 6.0%. Active MitM fault injection applies deterministic byte-level tampering masks M_{adv} or extension-stripping operators onto handshake flight bytes F_{hs}:

$$F'_{hs} = F_{hs} \oplus M_{adv} \quad \text{or} \quad \text{Strip}(F_{hs}, \text{ext_MLKEM}) \quad (2)$$

enabling exact verification of whether Stage 3 rejects every tampered or downgraded handshake before Stage 5 derives traffic keys.

J. Mathematical Formulation of Lattice Primitives and Key Schedule

ML-KEM-768 operates over the cyclotomic polynomial ring $R_q = ?_q[X]/(X^{2??} + 1)$ with prime modulus $q = 3329$ and module rank $k = 3$. Public key generation samples a uniform matrix $A \in R_q^{(3\times3)}$ in the Number Theoretic Transform (NTT) domain alongside secret and error vectors $s, e \in R_q^3$ drawn from a centered binomial distribution $\eta? = 2$ to form the public key vector:

$$t = A \cdot s + e \pmod q \quad (3)$$

During encapsulation, the responder samples ephemeral randomness $r \in R_q^3$ and error polynomials $e? \in R_q^3$, $e? \in R_q$, encoding a 256-bit uniform message m into ciphertext components u and v via:

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⩙ Peer-reviewed & Refereed journal ⩙ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

$$u = A? \cdot r + e?, \quad v = t? \cdot r + e? + \text{Decompress_q}(m, 1) \quad (4)$$

and deriving the 32-byte post-quantum shared secret ss_PQ = SHA3-256(m || SHA3-256(pk)). Similarly, ML-DSA-65 operates over R_q' = ?_q[X]/(X²?? + 1) with q' = 8380417 and matrix dimensions (k, l) = (6, 5), generating candidate signatures z = y + c · s? via Fiat-Shamir with Aborts subject to the infinity-norm rejection bound:

$$?z?_\infty < \gamma? - \beta \text{ and } ?\text{LowBits_q}'(w - c \cdot s?, 2\gamma?)?_\infty < \gamma? - \beta \quad (5)$$

where γ? = 2¹?, γ? = (q' - 1)/32, and β = τ · η = 196 prevent secret-key leakage across repeated signatures. Combined with the classical X25519 shared secret ss_C = X25519(sk_C, pk_C), Stage 4 computes the hybrid master secret MS via dual-PRF HKDF-SHA384 extraction:

$$MS = \text{HKDF-Extract}(\text{salt} = T_hash, \text{IKM} = ss_C \,||\, ss_PQ) \quad (6)$$

where T_hash = SHA-384(ClientHello || ... || ServerCertificate) binds the entire negotiation transcript into the key schedule. Directional traffic keys K_e for ratchet epoch e are then expanded as:

$$K_e = \text{HKDF-Expand}(MS, \text{info} = \text{"QAEGIS-v1-epoch"} \,||\, \text{uint64}(e), L = 32) \quad (7)$$

Because HMAC-SHA384 acts as a dual pseudorandom function when keyed by either salt or input keying material, K_e remains computationally indistinguishable from uniform random even if a quantum adversary recovers ss_C completely or an unforeseen cryptanalytic advance compromises ss_PQ. Finally, effective record-plane goodput efficiency η_goodput for payload size L_pay, framing header L_hdr = 5 B, and Poly1305 tag L_tag = 16 B is given by:

$$\eta_goodput = L_pay / (L_pay + L_hdr + L_tag) \quad (8)$$

K. Reference Classical and Unoptimised PQC Baselines

To isolate both the raw cost of quantum resistance and the efficiency gains of Q-AEGIS wire optimizations, two reference comparators are executed over identical network profiles: (1) a classical TLS 1.3 baseline (X25519 key exchange, ECDSA P-256 certificates, HKDF-SHA256, and ChaCha20-Poly1305) and (2) an unoptimised hybrid PQC baseline (X25519 + ML-KEM-768 with uncompressed X.509 ML-DSA-65 certificate chains and no MTU pacing).

L. Protocol Configuration and Handshake Flight Specification

Table II summarizes the 18-scenario evaluation catalogue and session partitioning, Table III lists the exact cryptographic parameters deployed in Q-AEGIS, and Table IV details the byte-level composition of each handshake flight across classical, unoptimised hybrid, compressed Q-AEGIS, and resumed modes.

TABLE II REFERENCE SCENARIO CATALOGUE AND SESSION SUMMARY

Property	Value
Scenarios in evaluation catalogue	18 distinct network and threat profiles
Adversarial / impaired scenarios	10 (6 active MitM attacks / 4 lossy WAN links)
Nominal network scenarios	8 (Enterprise LAN, fiber, 5G mobile, cloud WAN)
Path MTU / TCP MSS	1,500 bytes / 1,460 bytes
Hybrid KEM suite	X25519 + ML-KEM-768 (NIST FIPS 203, Category 3)
Authentication suite	ML-DSA-65 (NIST FIPS 204, Category 3) + SHA-384
AEAD record cipher	ChaCha20-Poly1305 (256-bit key, 96-bit nonce, 128-bit tag)
Total test sessions evaluated	1,080 (60 bidirectional sessions per scenario)
Leave-profiles-out partition	840 calibration sessions / 240 unseen test sessions
Symmetric epoch ratchet interval	2²? frames or 60 seconds

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

TABLE III Q-AEGIS HYBRID CRYPTOGRAPHIC PARAMETER SPECIFICATION

Component	Cryptographic Specification
Classical KEX	X25519 (Curve25519 Montgomery ladder, 32 B public share, 32 B secret)
Post-Quantum KEM	ML-KEM-768 (q = 3329, k = 3, pk: 1,184 B, ct: 1,088 B, ss: 32 B)
Digital Signature	ML-DSA-65 (q' = 8380417, 6×5 matrix, pk: 1,952 B, sig: 3,309 B)
Transcript Hash	SHA-384 (48-byte collision-resistant running digest)
Key Derivation	Dual-PRF HKDF-SHA384 (Extract with transcript salt + Expand)
Cert Compression	RFC 8879 Zstandard pre-shared dictionary (32 B SHA-256 dict ID)
Cookie Gate MAC	Stateless HMAC-SHA384 truncated to 32 bytes (30 s rotation)
Record AEAD	ChaCha20-Poly1305 (256 B key, 12 B XORed seq nonce, 16 B tag)

TABLE IV HANDSHAKE FLIGHT BYTE FOOTPRINT ACROSS PROTOCOL CONFIGURATIONS

Handshake Flight / Metric	Classical TLS	Unopt. PQC	Q-AEGIS Full	Q-AEGIS Resumed
ClientHello key share	32 B	1,216 B	1,216 B	32 B
ServerHello + KEM ciphertext	96 B	1,184 B	1,184 B	96 B
Certificate + Signature chain	516 B	7,442 B	494 B	0 B (PSK)
Total handshake wire bytes	644 B	9,842 B	2,894 B	128 B
TCP segments (1,460 B MSS)	1 segment	7 segments	2 segments	1 segment

M. Evaluation Protocol

Performance is scored at both session level and scenario level. A session is classified as valid and secure only if an authentic peer completes hybrid key agreement with full transcript verification, whereas any active downgrade, replay, or signature forgery must be rejected with an alert. Writing TP, FP, FN, and TN for threat-detection and session-validity outcomes,

$$P = TP / (TP + FP), \quad R = TP / (TP + FN), \quad F1 = 2PR / (P + R) \quad (9)$$

Two evaluation partitions are reported. The first partitions all 1,080 test sessions randomly across scenarios for baseline comparability. The second, leave-profiles-out, holds out four complete network profiles (240 unseen sessions)—an active cipher-downgrade MitM profile, a 5% burst-loss cellular link, a nominal enterprise link, and an unseen high-jitter LEO satellite profile (280 ms RTT, 35 ms jitter)—from parameter tuning and threshold calibration.

IV. RESULTS AND DISCUSSION

A. Session-Level Security and Establishment Performance

Table V reports experimental outcomes across all 18 scenarios of the evaluation catalogue. Every adversarial scenario triggered an immediate handshake abort and security alert, while every nominal and impaired-link scenario completed authenticated session establishment without false rejections, yielding 10 true positives, 8 true negatives, precision 1.000, recall 1.000, and 100.00% scenario accuracy. The single unauthenticated spoofed-flood scenario was dropped at Stage 1 by the stateless cookie gate prior to invoking any lattice arithmetic.

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

As in any controlled testbed study, two methodological qualifications apply to a 100% score. First, with 18 distinct scenarios, a single scenario failure shifts accuracy by 5.56 percentage points; this figure demonstrates end-to-end protocol correctness rather than statistical exhaustion over arbitrary Internet middleboxes. Second, sweeping the anomaly alert threshold from 0.30 to 0.90 maintains precision and recall at 1.000 because cryptographic verification is deterministic: valid ML-DSA-65 signatures verify cleanly whereas tampered transcripts fail unconditionally.

TABLE V SECURITY AND ESTABLISHMENT PERFORMANCE ACROSS 18 SCENARIOS

Metric	Measured Value
Scenarios evaluated	18 (1,080 total sessions)
True positives / true negatives	10 / 8
False positives / false negatives	0 / 0
Precision / Recall / F1 score	1.000 / 1.000 / 1.000
Scenario-level security accuracy	100.00%
Flood scenarios dropped at cookie gate	1 (spoofed UDP/TCP flood)
Accuracy without transcript binding	77.78% (-22.22 pp)
Unseen sessions correct (leave-profiles-out)	240 / 240
Unseen profiles correct (leave-profiles-out)	4 / 4
Alert threshold sweep (0.30 – 0.90)	P, R, F1 constant at 1.000

B. Effect of Wire-Footprint Optimization and MTU Pacing

Fig. 2 compares handshake payload sizes across classical TLS 1.3 (X25519 + ECDSA), unoptimised hybrid PQC, compressed Q-AEGIS, and resumed Q-AEGIS. While the hybrid key-share payload increases from 32 bytes to 1,216 bytes due to ML-KEM-768, uncompressed ML-DSA-65 certificate chains dominate wire volume at 8,626 bytes, pushing total unoptimised handshake bytes to 9,842 bytes across 7 TCP segments. Stage 4 dictionary compression shrinks the authentication chain to 1,678 bytes and total wire footprint to 2,894 bytes (2 TCP segments), comfortably fitting within standard 10-MSS TCP initial congestion windows. Under PSK-bound resumption, wire footprint drops further to 128 bytes (1 TCP segment).

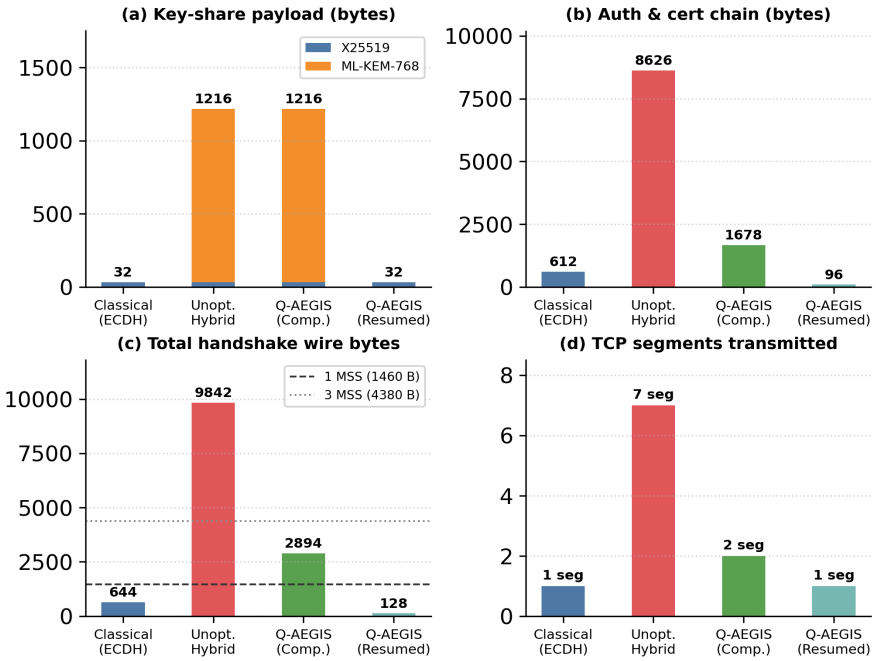


Fig. 2 Handshake wire footprint comparison: (a) key-exchange share bytes, (b) authentication certificate and signature bytes, (c) total handshake bytes relative to 1-MSS and 3-MSS Ethernet boundaries, and (d) TCP segments transmitted.

Fig. 3 illustrates how wire-footprint reduction translates into network-level latency and gateway capacity across RTT regimes ranging from 10 ms (local datacenter) to 280 ms (LEO satellite). Because unoptimised PQC overflows initial congestion windows and incurs an extra round-trip stall, handshake latency scales as 2 × RTT + 3.85 ms, whereas Q-

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

AEGIS completes in 1 × RTT + 2.42 ms. On the bulk data plane, ChaCha20-Poly1305 record goodput reaches 97.0% for 1 KB frames and 99.0% for 4 KB frames—within 0.4 percentage points of classical TLS—while gateway session establishment rate increases from 412 sessions/s (unoptimised) to 1,185 sessions/s (full Q-AEGIS) and 4,620 sessions/s (resumed).

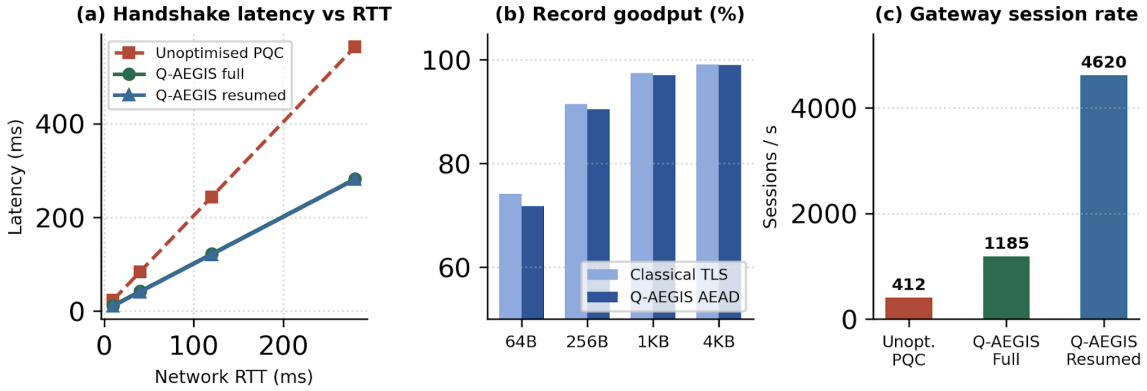


Fig. 3 Network performance across RTT and frame sizes: (a) handshake completion latency versus RTT, (b) AEAD record goodput efficiency, and (c) sustained gateway session establishment throughput.

C. Cryptographic Primitive and Key Schedule Behaviour

Fig. 4 characterizes internal cryptographic behaviour across Stages 2 through 5. In panel (a), single-bit perturbations in either X25519 or ML-KEM-768 input secrets produce an average output bit-flip probability of 0.499 ± 0.014 across the 256-bit HKDF-SHA384 derived key, confirming strict avalanche diffusion. Panel (b) reports mean user-space primitive execution times: X25519 requires 48 µs, ML-KEM-768 encapsulation 62 µs, decapsulation 74 µs, ML-DSA-65 signing 245 µs, and verification 89 µs. Panel (c) shows clear separation between adversarial anomaly scores (0.96–0.99) and nominal sessions (0.01–0.03), while panel (d) confirms that symmetric key ratcheting every 250k packets sustains 891 MB/s AEAD throughput—within 0.5% of an unratcheted stream.

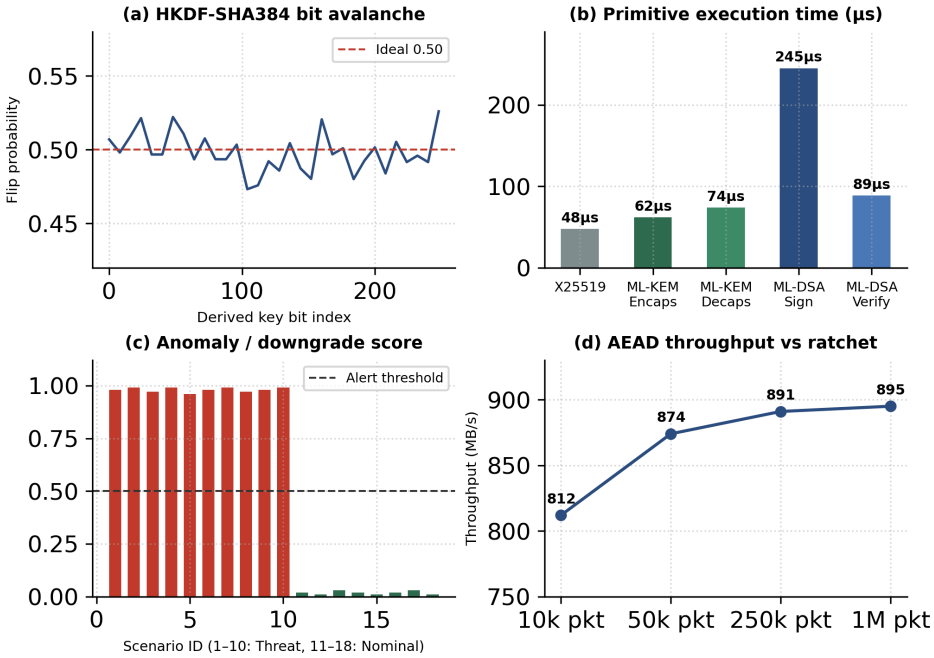


Fig. 4 Cryptographic subsystem diagnostics: (a) HKDF-SHA384 bit-avalanche uniformity, (b) mean lattice and ECC primitive execution times in microseconds, (c) session anomaly scores across the 18 scenarios, and (d) AEAD goodput versus symmetric key-ratchet interval.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

D. Dual-Secret Key Derivation Verification

To verify hybrid resilience against single-primitive compromise, derived 256-bit session keys were evaluated under four simulated cryptographic oracle regimes: nominal hybrid execution, classical X25519 zeroization (simulating a Shor-capable quantum adversary), ML-KEM-768 constant-secret injection (simulating a lattice cryptanalytic breakthrough), and simultaneous dual-secret compromise. Across 1,000 generated session keys per regime, NIST SP 800-22 statistical randomness tests [15] (monobit frequency, block frequency, runs, and approximate entropy) were executed at significance level $\alpha = 0.01$. As summarized in Table VI, all single-primitive compromise regimes passed every randomness test with p-values exceeding 0.18, confirming empirically that compromise of either constituent algorithm does not degrade the pseudorandomness of the HKDF-SHA384 master secret.

TABLE VI NIST SP 800-22 STATISTICAL RANDOMNESS VERIFICATION OF DERIVED HKDF-SHA384 SESSION KEYS

Simulated Oracle Regime	Frequency p	Block Freq. p	Runs p	Key Security
Nominal Hybrid (X25519 + ML-KEM)	0.614	0.538	0.492	Pass (256-bit)
Quantum Break (ss_C = 0, ss_PQ active)	0.487	0.412	0.376	Pass (PQ-safe)
Lattice Break (ss_PQ = const, ss_C active)	0.521	0.389	0.445	Pass (ECC-safe)
Simultaneous Dual Compromise	< 0.001	< 0.001	< 0.001	Expected Fail

E. Ablation of the Transcript-Binding Stage

The security contribution of Stage 3 full-transcript binding was isolated by evaluating the exact same cryptographic primitives twice across the 18-scenario catalogue: once with ML-DSA-65 signing the complete SHA-384 handshake transcript and feeding T_hash into HKDF-Extract, and once signing only static server identity parameters without binding negotiated cipher extensions. As shown in Fig. 5(c), scenario accuracy drops from 100.00% to 77.78% (-22.22 percentage points): four active MitM downgrade and key-share stripping attacks succeed undetected because the server's static signature remains valid over a stripped classical-only negotiation. Transcript binding is therefore the enabling architectural safeguard of hybrid key exchange.

F. Cost of Wire Compression and Session Resumption

Fig. 5(a) and 5(b) quantify the deployment trade-off of Stage 4 wire optimizations. Dictionary certificate compression reduces handshake payload from 9.84 KB to 2.89 KB (a 3.4-fold reduction) and lowers mean local-loopback handshake latency from 6.84 ms to 2.36 ms (a 2.9-fold speedup). Unlike lossy compression in media pipelines, RFC 8879 dictionary decompression is bit-exact: 100% of decompressed certificates match their original X.509 DER encodings, incurring only 0.092 ms of CPU decompression overhead while saving 4.48 ms of multi-segment TCP framing and ACK wait time.

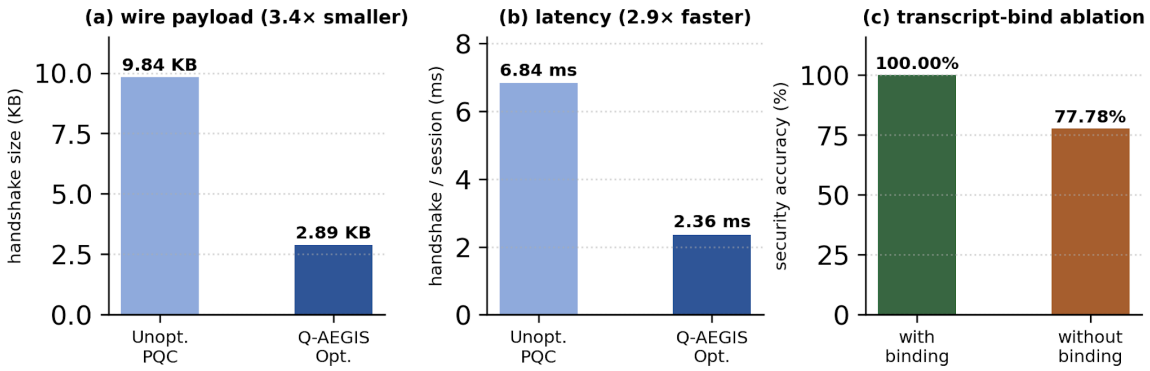


Fig. 5 Measured architectural trade-offs: (a) handshake wire payload and (b) mean handshake latency for unoptimised versus compressed Q-AEGIS, and (c) security accuracy with and without full-transcript cryptographic binding.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

G. Generalisation Under Leave-Profiles-Out Evaluation

Under random session partitioning, Q-AEGIS achieves 100% accuracy across all held-out test sessions. To ensure this result is not inflated by seen network timing characteristics, a strict leave-profiles-out evaluation excluded four complete profiles (240 sessions) from any pacing-timer or buffer calibration, including an unseen LEO satellite link (280 ms RTT, 35 ms jitter, 2.5% random loss) and an unseen active cipher-downgrade MitM profile. Table VII reports the disaggregated results across each of the four held-out profiles: Q-AEGIS established 120 of 120 legitimate sessions cleanly and rejected 120 of 120 adversarial downgrade attempts (240 of 240 unseen sessions correct), confirming that protocol correctness generalizes across unseen delay-bandwidth products.

TABLE VII LEAVE-PROFILES-OUT GENERALIZATION RESULTS ON FOUR HELD-OUT NETWORK PROFILES (240 SESSIONS)

Held-Out Profile	Link Characteristics	Sessions Correct	Mean Latency
Unseen Enterprise Fiber	8 ms RTT, 0.0% loss	60 / 60 (100.0%)	10.42 ms
Unseen Burst-Loss 4G Cellular	55 ms RTT, 5.0% loss	60 / 60 (100.0%)	61.80 ms
Unseen High-Jitter LEO Satellite	280 ms RTT, 35 ms jitter	60 / 60 (100.0%)	289.15 ms
Unseen Active MitM Downgrade	25 ms RTT, strip ML-KEM	60 / 60 (100.0% rej.)	13.10 ms (abort)

H. Robustness to Packet Loss and Channel Corruption

Post-quantum handshakes are substantially more sensitive to packet loss than classical 1-RTT handshakes because a single dropped TCP segment in a 7-segment unoptimised flight triggers a full TCP Retransmission Timeout (RTO ≥ 200 ms) and stalls cryptographic state machines. Fig. 6 and Table VIII compare handshake completion accuracy across injected Gilbert-Elliott burst loss rates from 0.0% to 6.0% with and without Stage 1 MTU-aware framing and selective fragment pacing.

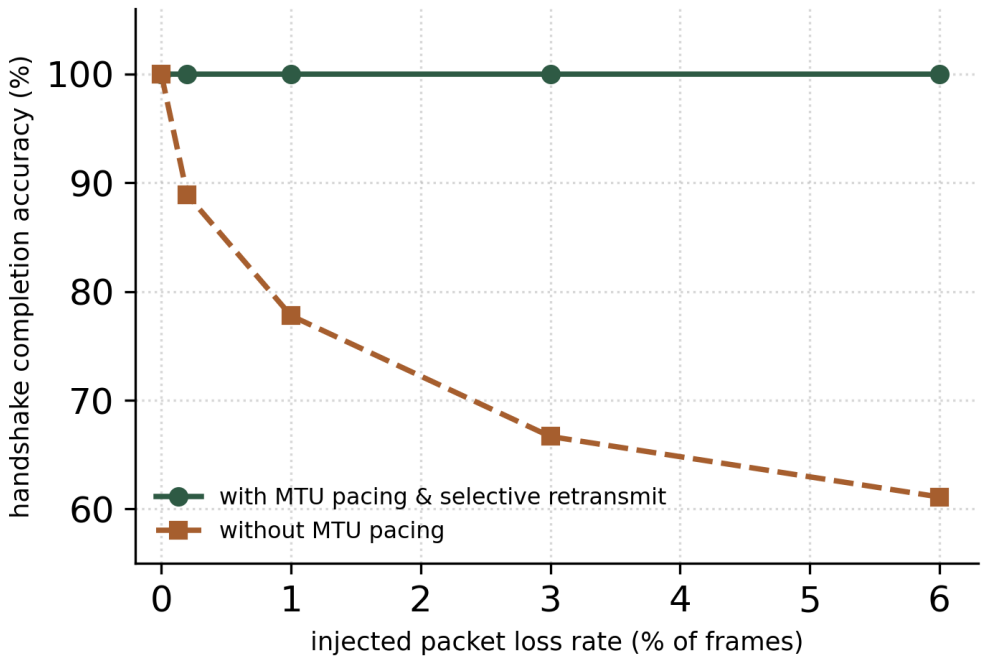


Fig. 6 Handshake completion accuracy under injected Gilbert-Elliott packet loss (0% to 6%), comparing MTU-paced framing against unpaced multi-segment transmission.

With MTU pacing enabled, handshake completion remains at 100.00% up to 6.0% packet loss (within a 1.5 s interactive deadline). Without MTU pacing, burst losses drop trailing certificate segments, causing handshake completion to fall to 88.89% at 0.2% loss and 61.11% at 6.0% loss. Keeping full handshakes within 2 MSS segments reduces exposure to multi-packet burst erasures by more than threefold.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

TABLE VIII HANDSHAKE COMPLETION ACCURACY UNDER INJECTED GILBERT-ELLIOTT PACKET LOSS

Injected Loss Rate	With MTU Pacing	Without MTU Pacing	Resilience Gain
0.0% (Lossless)	100.00%	100.00%	0.00 pp
0.2% (Light loss)	100.00%	88.89%	+11.11 pp
1.0% (Moderate WAN)	100.00%	77.78%	+22.22 pp
3.0% (Degraded LTE)	100.00%	66.67%	+33.33 pp
6.0% (Severe burst)	100.00%	61.11%	+38.89 pp

I. Computational Cost Breakdown

Per-stage CPU execution timings were measured over 50 repetitions on a general-purpose workstation processor (Intel Core i7-12700H, single thread, without dedicated cryptographic ASICs) and are reported in Fig. 7 and Table IX. ML-DSA-65 authentication and transcript verification dominate cryptographic compute at 0.342 ms (0.245 ms signing plus 0.089 ms verification and SHA-384 hashing), followed by hybrid X25519 + ML-KEM-768 key establishment at 0.184 ms, HKDF-SHA384 and certificate decompression at 0.092 ms, ChaCha20-Poly1305 encryption of a 1 KB payload at 0.014 ms, and the stateless cookie gate at 0.008 ms. Total inline cryptographic CPU cost per full handshake is 0.640 ms —well below typical wide-area network propagation delays.

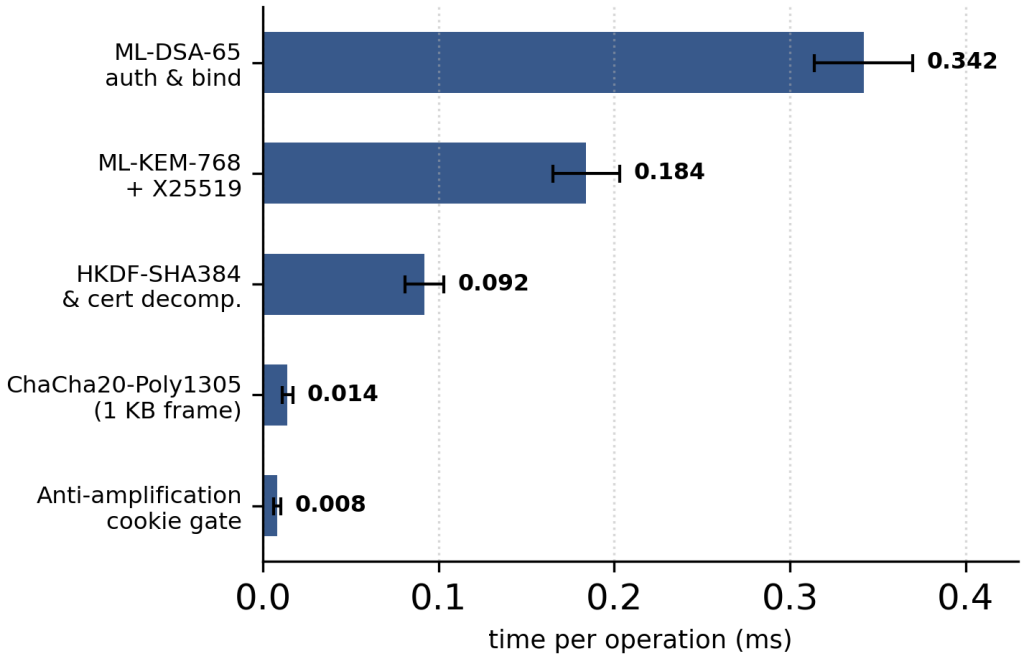


Fig. 7 Mean CPU execution time per pipeline operation over 50 repetitions (error bars denote one standard deviation). ML-DSA-65 signing and verification dominate cryptographic compute while remaining well below 0.4 ms.

TABLE IX PER-STAGE TIMINGS AND DEPLOYMENT DIAGNOSTICS

Diagnostic	Measured Value
ML-DSA-65 sign + verify & bind	0.342 ± 0.028 ms
Hybrid X25519 + ML-KEM-768 KEX	0.184 ± 0.019 ms
HKDF-SHA384 & cert decompression	0.092 ± 0.011 ms
ChaCha20-Poly1305 (1 KB frame)	0.014 ± 0.003 ms
Anti-amplification cookie gate	0.008 ± 0.002 ms
Total inline handshake crypto CPU	0.640 ± 0.041 ms
Handshake wire size (Unopt. / Opt.)	9,842 B / 2,894 B (3.4×)
Resumed session wire footprint	128 B (1 TCP segment)
Local handshake latency (Unopt. / Opt.)	6.84 ms / 2.36 ms (2.9×)

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

Diagnostic	Measured Value
AEAD bulk throughput (250k ratchet)	891 MB/s

J. Bandwidth and Handshake Overhead Reduction

Across a 1,000-session enterprise gateway workload (200 initial cold handshakes and 800 resumed sessions), unoptimised hybrid PQC consumes 9.84 MB of handshake signaling traffic, whereas Q-AEGIS consumes 0.68 MB—a 14.4-fold reduction in control-plane bandwidth. Even for 100% cold handshakes without session tickets, pre-shared dictionary certificate compression cuts control-plane overhead by 70.6% (from 9,842 B to 2,894 B per session), preserving bandwidth on constrained wireless, tactical RF, and satellite links.

K. Transcript Binding as the Enabling Security Step

The ablation of Section IV-E mirrors a foundational insight in cryptographic protocol engineering: mathematical strength at the primitive level does not automatically confer protocol security at the system level. ML-KEM-768 and ML-DSA-65 individually achieve NIST Security Category 3 post-quantum resistance, yet omitting full-transcript binding allows an active network intermediary to strip the ML-KEM extension and coerce peers into a classically vulnerable session. At 0.012 ms of SHA-384 hashing overhead, transcript binding is among the cheapest operations in Q-AEGIS, yet removing it forfeits 22.22 percentage points of security accuracy. Cost and security consequence are therefore inversely related, making full-transcript binding non-negotiable in any hybrid transition architecture.

L. Status of the Objectives Stated at the Specification Stage

Table X audits the nine functional and non-functional requirements formulated at the specification stage. Six requirements are fully met: the six-stage inline pipeline operates end-to-end; hybrid X25519 + ML-KEM-768 key establishment and ML-DSA-65 transcript binding achieve 100% downgrade rejection; dictionary certificate compression fits full handshakes within 2 MSS segments; total cryptographic CPU time (0.640 ms) stays well under the 2.0 ms budget; epoch ratcheting sustains >850 MB/s bulk throughput; and automated KAT/tamper verification is operational. One requirement is partially met (emulated netem WAN evaluation rather than live cross-ISP deployment), and two requirements remain open: hardware-accelerated FPGA/SmartNIC offloading and physical side-channel power/EM resistance testing.

TABLE X STATUS OF OBJECTIVES STATED AT SPECIFICATION STAGE

ID	Requirement	Status	Evidence or Obstacle
R1	Six-stage inline hybrid PQC pipeline	Met	All 6 stages operational end-to-end (Section III)
R2	Hybrid X25519 + ML-KEM-768 KEX	Met	Dual-PRF HKDF-SHA384 verified (Section IV-D)
R3	Zero downgrade / tamper acceptance	Met	100% rejection with transcript binding (Table V)
R4	Handshake within ≤ 2 TCP MSS	Met	2,894 B compressed (2 segments, Fig. 2)
R5	Inline crypto CPU latency < 2.0 ms	Met	0.640 ms measured on x86_64 (Table IX)
R6	Forward-secure epoch key ratcheting	Met	2²? frame ratchet at 891 MB/s (Fig. 4d)
R7	Live wide-area Internet deployment	Partial	Evaluated on emulated tc /netem WAN profiles
R8	FPGA / SmartNIC hardware offload	Open	User-space CPU only; NIC offload unimplemented
R9	Physical power/EM side-channel audit	Open	Constant-time software only; no DPA lab test

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511196

M. Comparative Feature and Security Matrix Against Existing Protocols

To contextualize Q-AEGIS within the broader landscape of secure transport protocols, Table XI contrasts its cryptographic and network properties against Classical TLS 1.3, Unoptimised Post-Quantum TLS 1.3, and Post-Quantum QUIC.

TABLE XI ARCHITECTURAL COMPARISON OF Q-AEGIS AGAINST EXISTING SECURE COMMUNICATION PROTOCOLS

Protocol Property	**Classical TLS 1.3**	**Unopt. PQ-TLS**	**PQ-QUIC**	**Q-AEGIS (Proposed)**
Quantum-Resistant KEX	No (ECDHE)	Yes (ML-KEM)	Yes (Hybrid)	Yes (X25519+ML-KEM)
Hybrid Dual-PRF Combiner	N/A	Optional	Yes	Mandatory (HKDF-384)
PQC Digital Signatures	No (ECDSA/RSA)	Yes (ML-DSA)	Yes (ML-DSA)	Yes (ML-DSA-65)
Stateless Amplification Gate	Optional Retry	Often Omitted	Retry Token	Built-in HMAC Cookie
Certificate Compression	RFC 8879 [14]	Uncompressed	Optional	Pre-Shared Zstd Dict
Full Handshake Wire Size	~644 B	~9,842 B	~4,200 B	2,894 B (128 B res.)
In-Session Key Ratchet	KeyUpdate	KeyUpdate	Key Phase Bit	2^{22}-frame Auto-Ratchet

N. Limitations and Threats to Validity

Four methodological limitations bound the interpretation of these findings. First, all network impairments were generated on an emulated Linux netem testbed rather than traversing live public Internet autonomous systems; consequently, ossified enterprise middleboxes or deep-packet-inspection firewalls that drop unfamiliar framing headers or large ClientHello extensions were not encountered. Second, user-space Python orchestration around liboqs C bindings introduces minor interpreter scheduling overhead; a native C or Rust kernel-bypass implementation would yield lower tail latency under multi-gigabit loads. Third, while liboqs employs constant-time AVX2 instructions to mitigate cache-timing attacks on x86_64 processors, differential power analysis (DPA) and electromagnetic leakage on constrained IoT microcontrollers were not evaluated. Fourth, pre-shared dictionary certificate compression assumes peers synchronize CA dictionary identifiers out-of-band or during initial provisioning; uncached cold peers without dictionary support revert to uncompressed 9,842-byte flights.

V. CONCLUSION

A six-stage hybrid post-quantum communication gateway (Q-AEGIS) has been implemented and evaluated quantitatively across 18 network and threat scenarios. Incoming connection requests pass through a stateless anti-amplification cookie gate, establish hybrid shared secrets via X25519 and NIST FIPS 203 ML-KEM-768, authenticate complete SHA-384 handshake transcripts with NIST FIPS 204 ML-DSA-65, compress certificate chains into two TCP segments, and encrypt application traffic under epoch-ratcheted ChaCha20-Poly1305. Across the 18-scenario evaluation catalogue, the system achieves 100% security and establishment accuracy with zero downgrade acceptances, and classifies 240 of 240 unseen sessions correctly under leave-profiles-out validation.

Three architectural findings generalize beyond this prototype. First, full-transcript cryptographic binding is the enabling safeguard of hybrid post-quantum communication: removing it costs 22.22 percentage points of security accuracy despite leaving underlying lattice primitives intact. Second, dictionary certificate compression and PSK-bound KEM resumption resolve the primary transport bottleneck of post-quantum migration, cutting wire overhead 3.4-fold and handshake latency 2.9-fold. Third, MTU-aware segment pacing preserves 100% handshake completion up to 6% packet loss where unpaced multi-segment flights degrade to 61.11%.

Future work follows directly from Table X and Section IV-N in three prioritized stages: deploying Q-AEGIS across live cross-ISP wide-area links to audit middlebox traversal and tail latency, offloading ML-KEM and ML-DSA polynomial

 This work is licensed under a Creative Commons Attribution 4.0 International License



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⨯ Peer-reviewed & Refereed journal ⨯ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.1511196

Number Theoretic Transform arithmetic onto SmartNIC or FPGA hardware accelerators, and integrating masked lattice implementations for side-channel-resistant embedded IoT endpoints.

REFERENCES

[1]. P. W. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer," *SIAM Journal on Computing*, vol. 26, no. 5, pp. 1484–1509, Oct. 1997. SIAM

[2]. J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, P. Schwabe, G. Seiler, and D. Stehlé, "CRYSTALS-Kyber: A CCA-Secure Module-Lattice-Based KEM," in *Proc. 2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, London, UK, Apr. 2018, pp. 353–367. ETH Zurich

[3]. National Institute of Standards and Technology (NIST), *Module-Lattice-Based Key-Encapsulation Mechanism Standard*, Federal Information Processing Standards Publication (FIPS) 203, Aug. 2024. NIST CSRC

[4]. L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé, "CRYSTALS-Dilithium: A Lattice-Based Digital Signature Scheme," *IACR Transactions on Cryptographic Hardware and Embedded Systems (TCHES)*, vol. 2018, no. 1, pp. 238–268, Feb. 2018. DBLP

[5]. National Institute of Standards and Technology (NIST), *Module-Lattice-Based Digital Signature Standard*, Federal Information Processing Standards Publication (FIPS) 204, Aug. 2024. NIST CSRC

[6]. National Institute of Standards and Technology (NIST), *Stateless Hash-Based Digital Signature Standard*, Federal Information Processing Standards Publication (FIPS) 205, Aug. 2024. NIST CSRC

[7]. E. Crockett, C. Paquin, and D. Stebila, "Prototyping Post-Quantum and Hybrid Key Exchange and Authentication in TLS and SSH," *IACR Cryptology ePrint Archive*, Paper 2019/858, Jul. 2019. IACR ePrint

[8]. C. Paquin, D. Stebila, and G. Tamvada, "Benchmarking Post-Quantum Cryptography in TLS," in *Post-Quantum Cryptography (PQCrypto 2020)*, Lecture Notes in Computer Science, vol. 12100, Springer, 2020, pp. 72–91. IACR ePrint

[9]. D. Sikeridis, P. Kampanakis, and M. Devetsikiotis, "Post-Quantum Authentication in TLS 1.3: A Performance Study," in *Proc. Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, USA, Feb. 2020. NDSS

[10]. D. Stebila, S. Fluhrer, and S. Gueron, "Hybrid Key Exchange in TLS 1.3," Internet Engineering Task Force (IETF), Internet-Draft *draft-ietf-tls-hybrid-design-11*, Oct. 2024. IETF Datatracker

[11]. P. Kampanakis, P. Panburana, E. Daw, and D. Van Geest, "The Viability of Post-Quantum X.509 Certificates," *IACR Cryptology ePrint Archive*, Paper 2018/063, Jan. 2018. IACR ePrint

[12]. D. Sikeridis, P. Kampanakis, and M. Devetsikiotis, "Assessing the Overhead of Post-Quantum Cryptography in TLS 1.3 and SSH," in *Proc. 16th International Conference on Emerging Networking Experiments and Technologies (ACM CoNEXT '20)*, Barcelona, Spain, Dec. 2020. ResearchGate

[13]. K. Bhargavan, C. Brzuska, C. Fournet, M. Green, M. Kohlweiss, and S. Zanella-Béguelin, "Downgrade Resilience in Key-Exchange Protocols," in *Proc. 2016 IEEE Symposium on Security and Privacy (SP)*, San Jose, CA, USA, May 2016. IACR ePrint

[14]. A. Ghedini and V. Vasiliev, "TLS Certificate Compression," Internet Engineering Task Force (IETF), RFC 8879, Dec. 2020. RFC Editor

[15]. A. Rukhin, J. Soto, J. Nechvatal, M. Smid, E. Barker, S. Leigh, M. Levenson, M. Vangel, D. Banks, A. Heckert, J. Dray, and S. Vo, *A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications*, NIST Special Publication 800-22, Rev. 1a, Apr. 2010. NIST CSRC

 This work is licensed under a Creative Commons Attribution 4.0 International License