

7th INTERNATIONAL CONFERENCE ON PUBLIC KEY INFRASTRUCTURE AND ITS APPLICATIONS (PKIA 2026)

SEPTEMBER 10-11th, 2026

**AH-KEM: Adaptive Hybrid Key Encapsulation Mechanism for
Post-Quantum Secure Cloud Communication**


Dvvit Patel • Krish Patel • Aaryan Lunis • Deepa Krishnan

Department of Computer Engineering, MPSTME, SVKM's NMIMS, Mumbai, India

Motivation: The Quantum Threat to Cloud Security

- RSA-2048 and ECDH secure today's TLS 1.3, SSH and cloud APIs — both broken in polynomial time by Shor's algorithm
- Harvest-Now, Decrypt-Later (HNDL): adversaries archive encrypted cloud traffic today to decrypt once a cryptographically-relevant quantum computer (CRQC) exists
- Mosca's inequality ($x + y > z$): if migration time exceeds the gap between data lifetime and CRQC arrival, confidentiality is already lost
- Long-lived records — health, financial, government — make today's classical handshake tomorrow's liability
- Migration is not “replace everything with PQC” — uniform deployment ignores cloud workload heterogeneity and wastes resources

MOSCA'S INEQUALITY



$x + y > z \Rightarrow$ ***data is exposed before migration finishes***

Motivates migrating ahead of any near-term CRQC timeline, not after one arrives.

Research Gap & Contributions


GAPS IN PRIOR WORK

- Homogeneous PQC deployment ignores heterogeneous cloud risk — every session gets the same cryptographic weight regardless of sensitivity
- No validated, adaptive hybrid-KEM combiner evaluated at scale for cloud environments
- Concurrency and scalability of PQC/hybrid handshakes under load remain unevaluated on consumer/cloud-grade hardware

AH-KEM CONTRIBUTIONS

- 1 Risk-based protocol with four cryptographic modes selected per session score
- 2 Formally-grounded Dual-PRF session-key combiner — IND-CCA2 secure in the QROM
- 3 Three-axis risk engine calibrated against NIST SP 800-57 key-management guidance
- 4 Benchmarks on Ryzen 7 7840HS: ML-KEM-768 keygen >11,000× faster than RSA-2048 (equivalent ephemeral setting)
- 5 Scalability validated: 0% errors to 250 concurrent clients, 100.0% risk-classification accuracy

AH-KEM System Architecture: Five-Message Handshake



Risk Engine scores the session and the Mode Dispatcher selects one of four modes (Table I); KEM Core runs ML-KEM/ML-DSA via the liboqs library; the Dual-PRF Combiner derives session key K_s independently on each side. (After Fig. 1)

Risk-Based Selection of Cryptographic Mode

$$R = 2s^d + 21T/60 + 3e_{xt}$$

$s^d \in \{1..5\}$
data sensitivity (policy-assigned)


T (min)
session duration, from ClientHello

$e_{xt} \in \{0,1\}$
external network exposure

Score	NIST Level	KEM(s)	Signature
0–3	Classical	X25519 only	ECDSA-P256
4–6	Hybrid L1	X25519 + ML-KEM-512	ML-DSA-44
7–9	PQC L3	ML-KEM-768 only	ML-DSA-65
10+	Hybrid L5	X25519 + ML-KEM-1024	ML-DSA-87

Lower risk → lighter classical protection. Higher risk → stronger PQC / hybrid protection. Weights calibrated to NIST SP 800-57 Part 1.

Dual-PRF Session-Key Combiner



$$\text{salt} = \text{SHA3-256}(K_c) \oplus \text{SHA3-256}(K_Q)$$

$$K_s = \text{HKDF-SHA3-256}(K_c \parallel K_Q, \text{salt}, \text{info})$$

Conservative hybrid security: K_s stays secure if at least one of X25519 or ML-KEM remains unbroken (QRom, Krawczyk's dual-PRF property).

FIVE-MESSAGE HANDSHAKE

- 1 ClientHello**
Ephemeral X25519 key + session metadata
- 2 ServerHello**
Risk score → mode M + KEM public keys
- 3 ClientKEM**
Encapsulation — ciphertext c_Q returned
- 4 ServerFinished**
Decapsulate, derive K_s, sign transcript (ML-DSA)
- 5 ClientVerify**
Verify signature, derive K_s, channel established

Implementation & Experimental Setup

LANGUAGE

Python 3.11

PQC LIBRARY

liboqs 0.10.0 (real ML-KEM / ML-DSA ops)

CLASSICAL KEM

X25519 — PyCA cryptography

CONCURRENCY

asyncio, non-blocking server

HARDWARE

Lenovo LOQ 15APH8 — AMD Ryzen 7 7840HS

OS / NETWORK

Native Linux, loopback (127.0.0.1)

1,000

iterations per mode
(latency benchmarking)

500

concurrent clients
(scalability test)

10,000

synthetic sessions
(risk-engine evaluation)

Performance: Key Generation & Handshake Latency

Mode / Baseline	Keygen (ms)	Encap (ms)	Sign (ms)	E2E (ms)
RSA-2048 + ECDSA	44.335	0.001	0.000	44.338
X25519 + ECDSA	0.061	0.114	0.000	0.177
Classical-only	0.063	0.116	0.000	0.180
Par.-Hybrid	0.065	0.175	0.129	0.384
PQC-Dominant	0.004	0.086	0.178	0.279
Full-Hybrid	0.065	0.214	0.226	0.514

1,000 iterations · Lenovo LOQ 15APH8 (Ryzen 7 7840HS) · 95% CI half-width < ±0.031 ms for all modes (Table II)

* RSA-2048 is measured with a fresh keypair generated per session (forward-secrecy setting), which deployed TLS does not normally do — the comparison isolates the cost of per-session ephemeral generation, not RSA's typical amortized cost.

>11,000×

faster ML-KEM-768 keygen vs RSA-2048*
(0.004 ms vs 44.335 ms)

0.514 ms


Full-Hybrid end-to-end handshake
(+0.337 ms over X25519 baseline)

0.279 ms

PQC-Dominant (NIST L3) median latency — well within cloud API SLA budgets


Scalability & Risk-Engine Accuracy

FULL-HYBRID SCALABILITY (Table IV)



0% errors up to 250 concurrent clients; 2.3% error at 500 clients as CPU saturates (94.1%).

RISK-ENGINE MODE DISTRIBUTION — n=10,000 (Table V)



100.0%


risk classification
accuracy

99.9%

of sessions routed to
quantum-resistant modes

Deployment Feasibility & Staged Migration

- Organizations on RSA/TLS see a net latency reduction moving to AH-KEM — removing the main performance objection to PQC adoption
- Cryptographic agility: sensitivity weights (w_s , w^T , w_e) live in a policy file, tunable without touching protocol code as standards evolve



Limitations & Conclusion

LIMITATIONS & FUTURE WORK

- Evaluated on a single local machine, not a geographically distributed deployment
- Native liboqs hardware acceleration (AVX-512) not yet exploited
- Static policy weights work well now — ML-based dynamic risk policies are future work
- Needs testing under realistic latency and packet-loss conditions
- Interoperability with existing TLS stacks and certificate infrastructure still to be verified

KEY FINDINGS

ML-KEM-768 keygen >11,000× faster than RSA-2048
under equivalent per-session ephemeral generation

Full-Hybrid handshake adds only 0.337 ms overhead
versus a classical X25519 baseline

100.0% risk-classification accuracy
across 10,000 synthetic sessions

99.9% of sessions routed to quantum-resistant modes
minimizing harvest-now-decrypt-later exposure

AH-KEM offers a practical adaptive bridge between classical and post-quantum cryptography — selecting protection by session risk while keeping measured latency low.

THANK YOU