International Journal of Applied Smart Interdisciplinary Technologies (IJASIT)

Home Page: https://ijasit.org **ISSN: 3139-759X**

https://doi.org/ 10.68104/ijasit.v1.i3.26

Hybrid Post-Quantum Cryptography and Quantum Key Distribution Co-Design for Quantum-Safe 6G, O-RAN, and IoT Infrastructure

Aagnik Chakraborty*

[1] Department of Computer Engineering, School of Engineering, P. P. Savani University (PPSU), Surat, Gujarat, India.
aagnikchakraborty@gmail.com

*** Corresponding author.**
E-mail: *aagnikchakraborty@gmail.com*

Article History:
Received Date: 07 August 2026
Revised Date: 18 August 2026
Accepted Date: 20 August 2026
Published Date: 10 September 2026

Keywords:

Post-quantum cryptography; quantum key distribution; Open RAN security; 6G infrastructure; cryptographic agility

ABSTRACT

Cryptographically relevant quantum computers do not yet exist, but the transition away from classical public-key cryptography is already underway: adversaries can record encrypted traffic today and decrypt it once a sufficiently capable quantum computer becomes available, a threat known as harvest-now, decrypt-later. This paper examines the engineering case for combining post-quantum cryptography (PQC) with quantum key distribution (QKD) as a strategy for securing next-generation communication infrastructure, focusing on 6G core and radio-access networks, Open RAN (O-RAN), and constrained Internet of Things (IoT) deployments. Rather than treating the two mechanisms as symmetric, interchangeable layers, the paper examines what each one actually guarantees, surveys the empirical evidence available for deployment cost, and describes where national security authorities and standards bodies disagree on the value of QKD. The architecture is grounded in the finalized National Institute of Standards and Technology (NIST) PQC standards and its 2025 hybrid key-establishment guidance, drawing on empirical Open RAN latency and energy measurements published during 2025 and 2026, embedded-hardware benchmarks for constrained devices, and the small number of quantum key distribution deployments operating at metropolitan or national backbone scale today. The paper is explicit about which figures are measured results and which remain illustrative proposals, and it closes by naming the specific standardization and field-trial gaps that separate current practice from the architecture it describes. The purpose of the work is to give network architects and policymakers a single, evidence-graded reference for deciding where each mechanism belongs in a real deployment, rather than treating post-quantum and quantum-distributed keying as equivalent options.

1. Introduction

Shor's algorithm, when run on a sufficiently large fault-tolerant quantum computer, solves the integer factorization and discrete logarithm problems in polynomial time [1]. Since RSA and elliptic-curve cryptography (ECC) derive their security from the classical intractability of exactly these problems, a cryptographically relevant quantum computer (CRQC) would break essentially all public-key cryptography deployed on the internet today, including TLS/HTTPS key exchange, code-signing, and the authentication primitives embedded throughout mobile-network standards.

No CRQC exists as of this writing. Expert timelines for when one might exist span a wide range and are genuinely uncertain; what is not uncertain is that an adversary can record ciphertext today and hold it for future decryption once a CRQC is available. This harvest-now, decrypt-later (HNDL) pattern is why national authorities — the U.S. National Security Agency (NSA), the UK National Cyber Security Centre, and NIST among them — have pushed migration timelines that begin years before any credible CRQC is expected, rather than waiting for one to appear.

1.1 Scope and Contribution of This Paper

This is a synthesis and architecture paper, not a report of new cryptanalysis or a new physical experiment. Its contribution is to assemble, in one place, (i) an architecture for combining classical, post-quantum, and quantum-distributed keying material without a single point of cryptographic failure; (ii) the actual empirical evidence — as of mid-2026 — for what that architecture costs in latency, energy, and memory across the telecom and IoT domains where it would need to run; and (iii) an honest account of where the U.S. and European positions on QKD diverge, because that disagreement is directly relevant to anyone deciding whether to fund QKD infrastructure alongside a PQC migration.

A note on evidentiary standard used throughout this paper: every quantitative figure is either (a) drawn from a NIST standard or a peer-reviewed/arXiv-preprint measurement study, cited at the point of use, or (b) explicitly labeled as an illustrative estimate, in which case the assumptions behind it are stated rather than presented as a measured result. Section 7 names, without euphemism, the specific evidence that does not yet exist.

2. Cryptographic Foundations

2.1 NIST's Finalized PQC Standards

In August 2024, NIST finalized the first three post-quantum cryptography standards to emerge from its 2016–2024 standardization competition [2, 3, 4]:

FIPS 203 — ML-KEM (Module-Lattice-Based Key-Encapsulation Mechanism, derived from CRYSTALS-Kyber): a KEM whose IND-CCA2 security rests on the hardness of the Module Learning-With-Errors (MLWE) problem over the ring $Zq[x]/(x256 + 1)$, $q = 3329$.

FIPS 204 — ML-DSA (Module-Lattice-Based Digital Signature Algorithm, derived from CRYSTALS-Dilithium): a Fiat–Shamir-with-aborts signature scheme built on MLWE and Module-SIS.

FIPS 205 — SLH-DSA (Stateless Hash-Based Digital Signature Algorithm, derived from SPHINCS+): a conservative, hash-only signature scheme offered as a structurally independent fallback in case a future weakness is found in lattice assumptions.

Table 1 gives the standardized parameter sizes. These numbers matter operationally: an ML-KEM public key is roughly 15–50× larger than an X25519 public key, and an ML-DSA-65 signature is on the order of 50× larger than a P-256 ECDSA signature. This is the concrete source of the packet-fragmentation, MTU, and TLS-certificate-chain-inflation problems that recur throughout the rest of this paper — not an abstract concern, but arithmetic that every constrained link budget in Sections 5 and 6 has to absorb. ML-KEM's shared secret is a fixed 32 bytes regardless of parameter set, which is what allows it to slot into existing symmetric key schedules (e.g. the TLS 1.3 key schedule) without protocol redesign.

Table 1 NIST-standardized PQC parameter sizes (bytes). Sources: FIPS 203 [2], FIPS 204 [3], FIPS 205 [4]

Standard	Parameter set	Security level	Public key	Secret key	Ct./Sig.
ML-KEM	ML-KEM-512	Cat. 1 (≈AES-128)	800	1,632	768 (ct.)

ML-KEM	ML-KEM-768	Cat. 3 (≈AES-192)	1,184	2,400	1,088 (ct.)
ML-KEM	ML-KEM-1024	Cat. 5 (≈AES-256)	1,568	3,168	1,568 (ct.)
ML-DSA	ML-DSA-44	Cat. 2	1,312	2,528	2,420 (sig.)
ML-DSA	ML-DSA-65	Cat. 3	1,952	4,000	≈3,300 (sig.)a
ML-DSA	ML-DSA-87	Cat. 5	2,592	4,864	≈4,600 (sig.)a
SLH-DSA	SLH-DSA-SHA2-128f	Cat. 1	32	64	17,088 (sig.)
SLH-DSA	SLH-DSA-SHA2-256f	Cat. 5	64	128	49,856 (sig.)

a Secondary sources report ML-DSA signature lengths that differ by a few bytes (e.g. 3,293 vs. 3,309 for ML-DSA-65) depending on encoding convention; we report an approximate figure rather than adjudicate the discrepancy.

2.2 Hybrid Key Establishment: NIST SP 800-227 and X-Wing

A pure PQC deployment is not, by itself, the conservative choice: ML-KEM and ML-DSA are seven and six years old respectively as public algorithms, versus multiple decades of cryptanalytic attention on RSA and ECC. The standards bodies' answer is hybrid key establishment: run a classical KEM/KEX and a post-quantum KEM in parallel and combine their outputs, so that an attacker must break both components to recover the session key.

NIST formalized this in SP 800-227, Recommendations for Key-Encapsulation Mechanisms, published in final form on 18 September 2025 after a public draft and workshop process [5]. Section 4.6 of SP 800-227 is devoted to multi-algorithm ("PQ/T hybrid") KEMs and gives X-Wing — the composite of ML-KEM-768 and X25519, since standardized by the IETF as RFC 9496 — as its worked example [5]. The design goal stated in the guidance is that the composite construction should remain secure as long as at least one of the components KEMs remains unbroken [5].

One detail from the final (as opposed to the initial public draft) version of SP 800-227 is worth flagging because it is the kind of implementation trap that looks harmless and is not: the final text tightened its rule on ephemeral key reuse (Requirement RS6) to explicitly forbid using an ephemeral KEM keypair for more than one key-establishment transaction, in direct response to public comment warning that ambiguity here had been read by some implementers as permitting short-lived reuse [5]. A hybrid combiner that is cryptographically sound on paper can still fail in exactly this way if the ephemeral-key discipline is not enforced in code.

2.3 The Combiner Construction and Why Domain Separation Matters

The generic dual-KEM combiner used throughout this paper's architecture is a keyed, domain-separated hash of the concatenated component secrets, given in Eq. (1):

$$k_{session} = KDF(\ k_{PQC} \parallel k_{classical} \parallel k_{QKD} \parallel ct_{PQC} \parallel ct_{classical} \parallel label\ ) \quad (1)$$

where KDF is an approved extract-and-expand function (e.g. HKDF instantiated with SHA-3), and label is a fixed, protocol-specific domain-separation string. The inclusion of the ciphertexts and a label is not decorative: a combiner that hashes only the raw shared secrets, without binding to the transcript and to a protocol identifier, is vulnerable to cross-protocol and key-reuse attacks in which a secret negotiated in one context is replayed to produce a colliding key in another. This is precisely why X-Wing's specification fixes a specific combiner and labeling scheme rather than leaving concatenation-and-hash as a free implementation choice [5]. Architectures that describe their combiner only as "H(k1‖k2)" without transcript binding are describing an insecure simplification; we state the full construction here so that Section 3's protocol integration points are unambiguous about what "combiner" means.

Because QKD-derived key material has no ciphertext to bind into the transcript in the same way, in practice kQKD is combined via a separate, authenticated channel-binding step (Section 4); we return to this when discussing PKI integration in Section 11.

3. A Layered Hybrid Architecture

3.1 Three-Layer Model

Figure 1 shows the architecture this paper argues for: classical asymmetric cryptography retained only for backward compatibility during the transition window, NIST-standardized PQC as the primary computational defense, and QKD reserved for the specific class of links where its trusted-node deployment model is economically and physically justified (Section 4) rather than deployed uniformly.

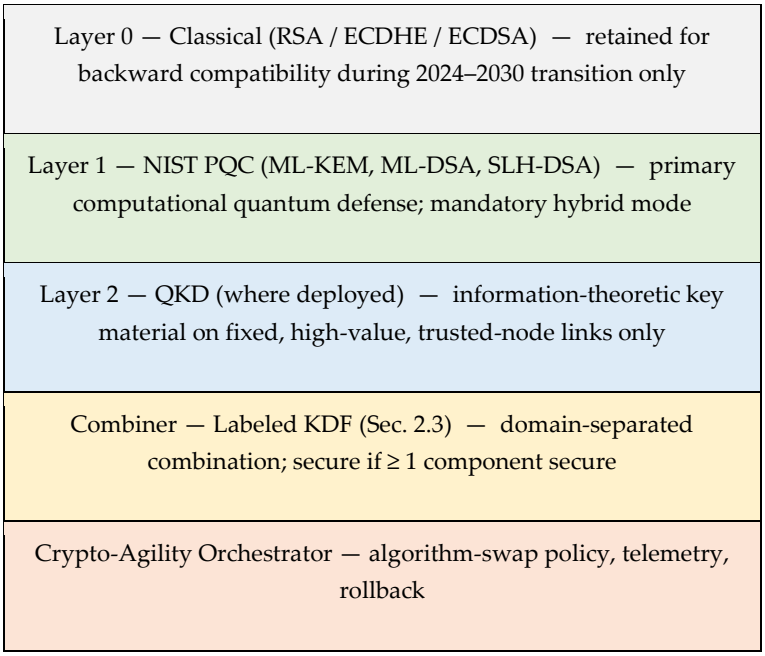

Fig. 1. Layered hybrid architecture. Only Layers 0–1 are present on every link; Layer 2 (QKD) is present only where the trusted-node deployment model is justified (Section 4).

3.2 Session Establishment and Fallback Logic

Session establishment proceeds in four phases: (1) a classical negotiation phase in which endpoints agree on algorithm sets; (2) a parallel hybrid key-exchange phase running the PQC KEM and, where present, the classical KEX and QKD key generation concurrently rather than sequentially — sequential composition, as mandated by earlier hybrid-IKEv2 drafts under RFC 9370, multiplies rather than adds latency penalties, which is one of the concrete reasons parallel composition is preferred in recent O-RAN hybrid designs [8]; (3) an authentication phase using ML-DSA-signed endpoint identities; and (4) derivation via the labeled combiner of Section 2.3.

Fallback is defined explicitly rather than left implicit, because an undefined fallback path is itself an attack surface (Section 8.2): if the QKD link is unavailable, the system falls back to PQC+classical hybrid; if PQC negotiation fails (e.g. due to a peer that has not yet upgraded), the system falls back to classical-only and this event is logged and alerted, since an attacker who can force this fallback on demand has defeated the hybrid design.

4. Quantum Key Distribution: Capability, Constraint, and the Standards Debate

4.1 What QKD Actually Provides

QKD protocols (BB84 and its decoy-state variants foremost among them) allow two parties who already share a short-authenticated channel to agree on a random key such that any eavesdropping attempt on the quantum channel introduces detectable disturbance, giving the key-exchange step information-theoretic (rather than computational) security. Two qualifications matter for systems engineering and are frequently elided in vendor material:

QKD secures only the key, not the data. The data itself is still encrypted with a conventional symmetric cipher (typically AES-256) keyed by the QKD output; QKD's information-theoretic guarantee does not extend to the confidentiality of the ciphertext against a classical adversary who compromises the endpoint or the symmetric cipher.

QKD requires an authenticated classical channel to bootstrap and to perform reconciliation and privacy amplification. That authentication step, today, is typically provided by a pre-shared key or — in a hybrid design — exactly the PQC/classical signature layer QKD is sometimes marketed as an alternative to. QKD does not eliminate the need for computational authentication; it complements it.

4.2 Deployment Reality: Distance, Trusted Nodes, and Scale

Direct fiber QKD without repeaters is limited by photon loss to roughly 100–300 km depending on detector technology; a 2026 field demonstration over deployed Swedish fiber using superconducting nanowire single-photon detectors (SNSPDs) reached 303 km with a single trusted-node relay [21]. Beyond that range, current deployments rely on trusted-node relays: intermediate sites that measure and re-transmit the key in the clear, meaning each relay node must itself be physically and operationally trusted — a real, non-cryptographic point of failure, and precisely the "single point of failure" concern this paper's overall architecture is meant to avoid.

China operates the world's largest deployed QKD infrastructure. The Beijing–Shanghai backbone, built by a USTC team led by Pan Jianwei and operational since the mid-2010s, spans roughly 2,032 km via 32 trusted relay nodes connecting Beijing, Jinan, Hefei, and Shanghai [18, 19]. This backbone has since been extended into the China Quantum Communication Network (CN-QCN), reported in npj Quantum Information in 2025 as spanning over 10,000 km with 145 backbone nodes and 20 metropolitan networks across 17 provinces and 80 cities, integrated with the Micius satellite for free-space links beyond fiber range [19]. Smaller metropolitan deployments exist elsewhere: a three-node, 25 km QKD network operated continuously in Cambridge, UK for 580 days, generating 120 Tb of secure key over its lifetime while encrypting two 100 Gbps classical channels [20]; the Madrid Quantum Communication Infrastructure (MadQCI) operates as a multi-vendor, multi-operator metropolitan testbed [20].

What does not yet exist, publicly, is a QKD deployment integrated end-to-end with an O-RAN control plane or an IoT gateway architecture at the scale this paper's 6G/O-RAN sections discuss. The deployments above are metropolitan or national backbone links — exactly the class of fixed, high-value, point-to-point connection where the trusted-node cost is justifiable. Section 7 returns to this gap explicitly.

4.3 The NSA Position, and Why a Serious Architecture Paper Has to Engage With It

Any paper proposing QKD as a co-equal pillar alongside PQC has an obligation to address the fact that the United States' National Security Agency, in the Commercial National Security Algorithm Suite 2.0 (CNSA 2.0) FAQ and its later addenda, explicitly does not approve QKD for use on National Security Systems and advises agencies against investing in or deploying QKD without direct consultation [22, 23]. The Department of War's internal PQC transition memo takes the same position [24]. NSA's stated reasoning is threefold and is worth stating plainly rather than paraphrasing away: QKD requires specialized, costly hardware with a distinct maintenance burden; it secures only the key, not the data, as noted above; and its distance and trusted-node constraints make it impractical for the kind of widespread, heterogeneous deployment a national network requires [23].

This is a genuine, unresolved disagreement between two credible standards bodies, not a settled question this paper can wave away: NIST's SP 800-227 treats QKD as one legitimate input to a multi-algorithm hybrid [5], while the CNSA 2.0 position is that PQC hybrids are sufficient and QKD's operational cost is not justified for NSS [22]. Meanwhile, the European Union has taken the opposite strategic bet, funding EuroQCI as a continent-scale QKD program (Section 9) precisely because the trusted-node model is judged acceptable for the specific set of government-to-government and cross-border financial links EuroQCI targets.

Our position, consistent with the architecture in Figure 1: QKD earns a place in a hybrid design specifically on fixed, high-value, point-to-point backbone links where the operator already controls or can vet every relay node — inter-datacenter links, national backbone segments, and cross-border government links being the clearest cases, and the JPMorgan Chase deployment in Section 7.1 being a live commercial example of exactly this pattern. It is not, on current evidence, a defensible requirement for the dynamically-routed, heterogeneous control-plane traffic of an O-RAN near-real-time RIC or for endpoint-to-gateway IoT links, where PQC-classical hybrids under crypto-agile management are the appropriate — and, per NSA's stated position for NSS, the officially preferred — mechanism.

5. O-RAN Integration: Evidence From the 2025–2026 Literature

5.1 The E2 Interface Latency Question: What Has Actually Been Measured

Open RAN's disaggregation of the DU, CU, and RIC functions across separately-operated components introduces new cryptographic pressure points, the most latency-sensitive of which is the E2 interface connecting the gNB to the Near-Real-Time RAN Intelligent Controller (Near-RT RIC), whose closed-loop xApp control decisions typically operate on sub-second timescales.

The relevant empirical evidence here is recent and specific. Perera et al. (2026) built an open-source O-RAN testbed — srsRAN, Open5GS, FlexRIC, and strongSwan with liboqs — and measured IKEv2/IPsec tunnel-establishment latency on the E2 interface under three configurations: no IPsec, classical ECDH-based IPsec, and ML-KEM-based IPsec [6]. Their finding, from repeated automated runs rather than a single measurement: ML-KEM integration adds approximately 3–5 ms to tunnel-establishment latency compared to classical ECDH, with no observed disruption to Near-RT RIC xApp control-loop stability or session-resumption behavior [6]. A separate 2025 study comparing ML-KEM against alternative post-quantum KEMs (BIKE, HQC) on O-RAN interfaces found BIKE achieving the lowest latency in their testbed, with ML-KEM competitive but showing more run-to-run variability than BIKE, and HQC showing higher jitter despite fast raw throughput [8] — a useful reminder that "PQC" is not a monolithic latency number and that the specific algorithm choice interacts with the specific interface's timing requirements.

Both figures are consistent with the widely-cited (but unpublished in peer-reviewed form) figure of "a few milliseconds" overhead that circulates in telecom PQC discussions; the value of the Perera et al. study is that it is a reproducible, open testbed result rather than a vendor claim.

5.2 Energy-Aware Scheduling: The Crypto Policy rApp / SOS xApp Architecture

The architecture most relevant to the energy cost of PQC handshakes at O-RAN scale is described by Gupta, Alimohammadi, Shojafar, Mi, and Bhutta in a 2026 paper proposing an O-RAN-aligned split between a Crypto Policy rApp residing in the Non-Real-Time RIC, which defines the strategic security envelope (PQC suite selection, rotation policy), and a Security Operations Scheduling (SOS) xApp in the Near-RT RIC, which converts that policy into tactical handshake timing and placement decisions — batching non-urgent handshakes, prioritizing session resumption, and selecting parameters to minimize joules per secure connection while respecting slice SLAs [7]. Cryptographic enforcement in their design remains at standards-compliant endpoints: MACsec at the O-DU/O-RU on the Open Fronthaul, and IPsec at the xhaul tunnel terminators [7].

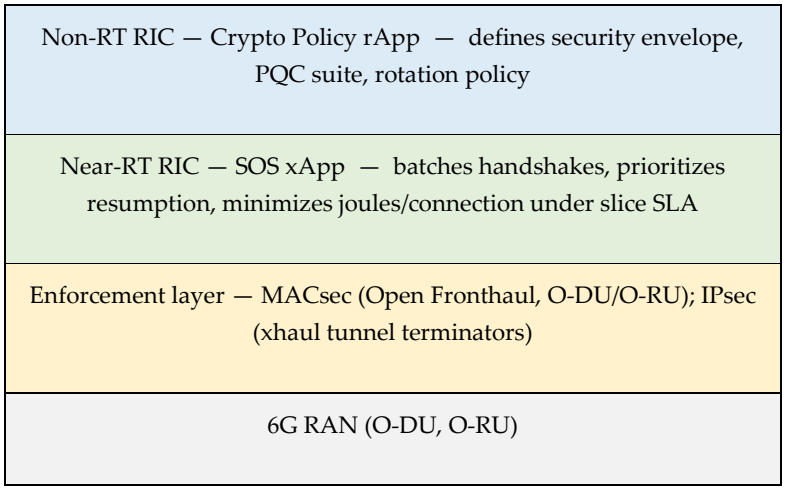

Fig. 2. The Crypto Policy rApp / SOS xApp energy-scheduling architecture, after Gupta et al. [7]. Evaluated via discrete-event simulation with literature-derived hardware benchmarks; not yet validated on physical O-RAN hardware.

We flag explicitly, because it matters for how this figure should be used: the reported result — that intelligent scheduling can reduce per-handshake energy by approximately 60% without violating slice latency targets — comes from a discrete-event simulation using 3GPP-aligned synthetic traffic profiles combined with hardware energy benchmarks drawn from the literature, not from a physical O-RAN deployment [7]. That is a legitimate and increasingly standard methodology for this kind of system-level evaluation, and the underlying

per-operation energy inputs are independently measured figures rather than assumptions (see Table 2), but the 60% headline figure describes a simulated system, and a physical validation of this specific scheduling architecture at O-RAN scale has not yet been published.

5.3 3GPP Standardization Status

3GPP's security working group SA3 initially deferred PQC work, stating in a 2020 liaison that it was "prudent to wait until NIST have completed their post-quantum cryptography project" before investigating quantum-safe algorithms for the 5G system [9]. That posture has since reversed: SA3's study item TR 33.703 now covers hybrid key exchange and PQC algorithm integration across RAN–core interfaces [10], and following 3GPP's Release 19 freeze in December 2025, SA3's PQC work has carried forward as part of the Release 20 study phase covering both 5G-Advanced and initial 6G work [11]. Ericsson, an active 3GPP contributor, has publicly stated its expectation that 5G standards will be updated to support PQC during the transition, and that 6G — targeted for deployment around 2030 — will carry full PQC support from its first release rather than as a bolt-on [12].

6. IoT at Scale: Constrained-Device Evidence

6.1 Device Tiering and the Delegation Pattern

The great majority of IoT endpoints cannot run PQC locally at all, let alone QKD. The standard architectural response — and the one this paper adopts — is to not make them: Class 0/1 devices (sensors, wearables, tens of KB of flash) perform only lightweight symmetric authentication (HMAC-SHA256) to a trusted edge gateway; the gateway, a Class 3 device with meaningfully more compute and power budget, performs the full PQC-hybrid negotiation with the cloud or telecom backend on the endpoint's behalf. This delegation pattern is not novel to this paper; it reflects the consensus architecture found throughout the IoT PQC literature, and its main engineering risk is that the gateway becomes a high-value single point of compromise, which is why gateway attestation and rapid credential rotation matter more here than raw handshake speed.

6.2 What Constrained Hardware Actually Measures

Table 2 reports real, independently reproducible benchmark figures rather than vendor estimates. Two results are worth calling out because they complicate the intuitive assumption that PQC is simply slower than what it replaces on constrained hardware:

On an STM32 Cortex-M4 using the pqm4-derived, assembly-optimized implementation, ML-KEM-512 (Kyber Level 1) completes a full handshake (keygen + encapsulation + decapsulation) in an average of 15.1 ms, versus 26.8 ms for ECDHE over SECP256R1 on the same hardware [13] — ML-KEM is faster, not slower, in this specific optimized comparison, because ECDH's point-multiplication cost on a general-purpose microcontroller is higher than ML-KEM's NTT-based polynomial arithmetic.

On the ARM Cortex-M0+ (Raspberry Pi Pico's RP2040, a genuinely resource-constrained target with no DSP extensions), a 2026 benchmarking study found ML-KEM-512 decapsulation comfortably fits within a 10 KB stack budget and the full handshake fits within the RP2040's 264 KB SRAM with room to spare, while ML-DSA signing is the more demanding operation, with ML-DSA-87 requiring roughly 120 KB of stack — about 45% of total available SRAM on that platform [15].

Table 2 Real, cited embedded-hardware PQC benchmarks (not vendor estimates)

Platform	Operation	Result	Source
STM32 (Cortex-M4)	ML-KEM-512 full handshake	15.1 ms avg. (vs. 26.8 ms for ECDHE/P-256)	[13]
STM32F4 (Cortex-M4, pqm4 reference)	ML-KEM-512 keygen (clean impl.)	≈595,800 cycles avg.	[14]
STM32F4 (Cortex-M4, pqm4 opt. impl.)	ML-KEM-768 keygen (m4fspeed)	≈392,400 cycles avg.	[14]
RP2040 (Cortex-M0+)	ML-KEM-512 decapsulation	9.4 KB peak stack; fits 10 KB budget	[15]
RP2040 (Cortex-M0+)	ML-DSA-87 signing	119.6 KB peak stack (≈45% of SRAM)	[15]

Raspberry Pi 4/5	ML-KEM-512 keygen throughput	3,807/s (Pi 4); 23,168/s (Pi 5)	[16]

The practical implication for the gateway-delegation architecture: an ML-KEM handshake is well within reach of even a Cortex-M0+-class gateway, but ML-DSA signature operations at higher security categories are stack-hungry enough that gateway firmware should budget for them explicitly rather than assume headroom, and endpoint devices below Class 2 should not attempt local PQC signing at all — a conclusion consistent with, and now given hard numbers by, the device-tiering pattern already standard in the field.

6.3 Side-Channel Considerations at the Endpoint

Lattice-based schemes' arithmetic — particularly non-constant-time polynomial sampling and message-encoding steps in early Kyber implementations — has been a documented target for power-analysis attacks; masked (randomized) implementations that eliminate secret-dependent branching and memory access have been published and analyzed specifically for Kyber on embedded targets [17]. This is real, published, ongoing work rather than a solved problem, and it is a legitimate reason for constrained gateway deployments to prefer vendor implementations that have undergone independent side-channel evaluation rather than assuming reference code is deployment-ready.

7. Deployment Evidence: What Exists Today

7.1 Financial Sector: JPMorgan Chase, Toshiba, and Ciena

The clearest real-world precedent for the fixed, high-value QKD link this paper argues for in Section 4 is JPMorgan Chase's collaboration with Toshiba and Ciena. In an initial 2022 demonstration, the partners multiplexed a QKD channel onto the same fiber as production 800 Gbps and 100 Gbps optical data channels over a 70 km link, at a key rate sufficient to support up to 258 AES-256-encrypted channels at a 1-key-per-second refresh rate, and used the resulting QKD-secured channel to protect Kinexys Liink, JPMorgan's production peer-to-peer blockchain network — described by the partners as the first demonstration of QKD securing a mission-critical blockchain application [25]. JPMorgan subsequently deployed a persistent Quantum-Secured Crypto-Agile Network (Q-CAN) connecting two data centers over a single 100 Gbps production fiber, explicitly pursuing what their Global CIO Lori Beer described as a dual remediation strategy combining PQC and QKD [26].

This is exactly the deployment pattern this paper's architecture recommends QKD for: a small number of fixed, high-value, operator-controlled fiber links between known endpoints, not a general-purpose control-plane mechanism.

7.2 National Backbone: China's QKD Network

China's Beijing–Shanghai backbone and its successor CN-QCN, described in Section 4, represent the only QKD deployment operating at national-backbone scale today, and its use case — connecting government facilities, banks, and (per CN-QCN reporting) grid-control infrastructure over fixed, state-vetted trusted-node relays [19] — is again consistent with the fixed-link recommendation above, not a counterexample to it.

7.3 The Honest Gap

We state this plainly because a paper that elides it is not credible to anyone who has tried to deploy this stack: as of mid-2026, there is no published, large-scale field deployment integrating QKD with an O-RAN near-real-time control plane, and no published field deployment of the specific Crypto Policy rApp / SOS xApp energy-scheduling architecture of Section 5 outside discrete-event simulation. The O-RAN latency evidence in Section 5 is real and testbed-measured; the energy-scheduling evidence is real and simulation-measured; the IoT hardware evidence in Section 6 is real and bench-measured; the QKD deployment evidence in this section is real and field-measured — but no single publication yet combines all four into one integrated, physically deployed system. That integration is the open engineering problem this paper's architecture proposes a design for, not a result it is reporting.

8. Threat Model and Crypto-Agility

8.1 Four Threat Classes

Store-now-decrypt-later (SNDL). Addressed directly by the hybrid design: an adversary must break both the PQC and classical components of a recorded session to recover it retroactively, and if QKD key material was used on that link, must also have compromised a trusted relay node at the time of transmission.

Cryptographic agility failure. If a weakness is found in a deployed PQC algorithm, the system needs a defined path to a secondary algorithm without a full redeployment cycle. NIST's own hybrid guidance is designed around this: because the combiner's security holds as long as one component is unbroken, a compromised PQC algorithm degrades the system to classical-only security rather than to zero security, buying time for an orchestrated swap [5]. We do not report a specific recovery-time figure here, because none has been measured in a realistic operational setting; the design property (graceful degradation, not sudden collapse) is what SP 800-227's combiner construction actually guarantees, and that is the property worth stating.

Side-channel exploitation. Discussed in Section 6; addressed by preferring independently-evaluated, masked implementations at the endpoint rather than reference code.

8.2 Downgrade and Fallback-Forcing Attacks

An attacker who can jam the QKD channel or block PQC negotiation — whether through active interference or by exploiting a misconfigured peer — can force a fallback to a weaker configuration. This is not a hypothetical concern specific to this architecture; downgrade attacks are a well-established class in protocol design generally. The mitigation is procedural as much as cryptographic: fallback events must be logged, rate-limited, and alerted on, with an explicitly defined minimum acceptable configuration (Section 3) rather than an implicit "do whatever negotiation succeeds" policy, since the latter is precisely what a downgrade attack exploits.

8.3 An Illustrative Availability Model

The following is offered as an illustrative modeling exercise, not a measured result, and we label it as such: if the three layers' compromise probabilities were independent with illustrative assumed values $P_{PQC} = 0.01$, $P_{classical} = 0.001$, $P_{QKD} = 0.001$ (chosen for illustration, not measured from any deployment), overall availability under the independence assumption would be given by Eq. (2):

$$A = (1 - P_{PQC})(1 - P_{classical})(1 - P_{QKD}) \approx 98.8\% \quad (2)$$

The value of this exercise is structural, not numerical: it shows that the hybrid design's benefit comes from the product of independent failure probabilities rather than their sum, which is the whole point of layering independent mechanisms — but the specific probabilities above are assumptions for illustration, and real component failure rates for any specific deployment would need to be measured, not assumed, before this model could inform an actual risk decision.

9. Policy and Regional Landscape

9.1 United States: NSM-10 and CNSA 2.0

National Security Memorandum 10 (May 2022) directs the U.S. federal government to migrate to quantum-resistant cryptography by 2035 [23]. CNSA 2.0, published by NSA in September 2022 and updated to version 2.1 in December 2024 following the FIPS 203/204/205 finalization, sets the concrete algorithm mandate for National Security Systems: ML-KEM-1024, ML-DSA-87, and (for firmware/software signing) LMS/XMSS or SLH-DSA-SHA2-256s, with a staged timeline — NSA ceased approving new RSA/ECC/Diffie-Hellman deployments for key establishment and signatures in 2025, hybrid classical+PQC solutions are expected through the 2025–2030 transition window, and NSA expects the substantial majority of NSS cryptography to have completed migration by 31 December 2031, with the full NSM-10 deadline at 2035 [23, 24]. As discussed in Section 4, this suite explicitly excludes QKD.

9.2 European Union: EuroQCI

The European Quantum Communication Infrastructure (EuroQCI) initiative, launched via a 2019 declaration now signed by all 27 EU member states, funds both national QKD network segments and cross-border interconnections through the Digital Europe Programme and the Connecting Europe Facility (CEF-Digital) [27]. Reported figures illustrate the initiative's real but still modest scale relative to broader digital-infrastructure spending: a 2025 CEF-Digital funding round allocated €389 million across 56 projects spanning submarine cable, 5G, and EuroQCI infrastructure jointly (not EuroQCI alone) [28]; a separate, EuroQCI-specific tracking figure from industry analysis reported approximately $97 million allocated to EuroQCI in 2024, roughly 8% of total EU quantum-technology investment that year [29]; and a dedicated €16 million call funds a QKD testing and certification facility running through 2027 [30].

The EU's approach is a direct strategic contrast to the U.S. CNSA 2.0 position: where NSA explicitly declines to fund QKD for NSS, the EU is funding QKD as core sovereign digital infrastructure, on the judgment that the

trusted-node model is acceptable for the specific government-to-government and cross-border financial links EuroQCI targets — the same class of fixed, high-value link this paper's Section 4 identifies as QKD's defensible use case.

9.3 China

China operates the CN-QCN backbone described in Section 4 and pursues a parallel track of domestically-developed cryptographic standards (including the SM2 elliptic-curve signature/encryption standard and the ZUC stream cipher used in mobile standards) alongside engagement with international PQC standardization; a detailed comparison of China's PQC-specific standardization timeline is outside the scope of what we can responsibly cite here without a dedicated review of Chinese-language standards documentation, and we do not attempt one.

10. Migration Roadmap

Rather than propose a new roadmap, we synthesize the timelines that the cited standards and vendor statements above actually commit to, since inventing a more precise one would misrepresent the state of consensus:

2024–2025: FIPS 203/204/205 finalized (August 2024); NSA ceases approving new RSA/ECC deployments for NSS key establishment and signatures (2025) [23]; SP 800-227 hybrid guidance finalized (September 2025) [5].

2025–2030: NSAs stated CNSA 2.0 transition window for hybrid classical+PQC solutions across NSS [23]; 3GPP SA3 continues PQC work into Release 20 following the Release 19 freeze (December 2025) [11]; EuroQCI national segments and cross-border links continue deployment under CEF-Digital [27].

~2030: Ericsson's publicly stated expectation of 6G deployment with full PQC support from the first release [12].

By 2031: NSA's expectation that the substantial majority of NSS cryptography has completed CNSA 2.0 migration [23].

By 2035: NSM-10's full quantum-resistant migration deadline for the U.S. federal government [23]. We deliberately do not add speculative dates (e.g. for O-RAN-integrated QKD field trials or a hypothetical TLS 1.4 hybrid standard) beyond what the cited sources commit to; Section 11 lists these as open questions rather than assigning them dates no standards body has yet committed to.

11. Open Research Questions

We close with the gaps that, in our reading of the current literature, most directly separate today's state of the art from the architecture proposed in Section 3:

1. QKD–PKI trust integration. There is no finalized protocol for how QKD-derived session keys should enter a certificate-based (PQC or classical) trust hierarchy in a way that is auditable and standardized across vendors; each deployment discussed in Section 7 currently uses a bespoke integration.

2. Hybrid combiner standardization beyond TLS 1.3. X-Wing/RFC 9496 addresses the two-component KEM case well; how a three-component PQC+classical+QKD combiner should be specified, labeled, and validated is not yet the subject of a finalized standard.

3. Physical validation of O-RAN energy-aware scheduling. The Crypto Policy rApp / SOS xApp architecture of Section 5 has strong simulation results [7] but, to our knowledge, no physical Near-RT RIC deployment result yet.

4. An integrated PQC+QKD+O-RAN or PQC+QKD+IoT field trial. As stated plainly in Section 7, this does not yet exist in the public literature; every component has been separately measured, but not the composed system.

5. Quantum repeater maturity. Extending QKD beyond the trusted-node model requires quantum memory-based repeaters capable of Bell-state measurement at practical rates; this remains a physics research problem, not an engineering integration problem, and timelines for it are genuinely uncertain rather than merely unannounced.

12. Conclusion

The engineering case for combining classical, post-quantum, and (selectively) quantum-distributed key material is sound, and it rests on real standards work: NIST's finalized FIPS 203/204/205 algorithms, its

September 2025 hybrid key-establishment guidance in SP 800-227, and a genuine, well-specified composite construction in X-Wing. The empirical cost of deploying PQC on the interfaces this paper examined is, on the evidence available, modest and specific: a few milliseconds of added latency on O-RAN's E2 interface [6], comfortably within reach of Cortex-M0+-class IoT gateways for KEM operations [15], and — per simulation, not yet physical measurement — amenable to roughly 60% energy reduction through intelligent scheduling [7]. Where this paper departs from treating PQC and QKD as symmetric, interchangeable pillars is the point we think matters most for anyone making a real deployment decision: NSA's National Security Systems guidance and the EU's EuroQCI program have made opposite strategic bets on QKD, for reasons grounded in real operational trade-offs rather than technical disagreement about the physics. A defensible architecture does not paper over that disagreement; it reserves QKD for the specific class of fixed, high-value, trusted-node-tolerant links — financial backbone, national backbone, government-to-government — where the JPMorgan Chase and China deployments described in Section 7 show the model already works commercially and at national scale, while relying on PQC-classical hybrids, managed under an explicit crypto-agility policy, for the dynamically-routed, heterogeneous control-plane and endpoint traffic that constitutes the overwhelming majority of 6G, O-RAN, and IoT communication.

References

[1]. Shor, Peter W. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer." SIAM Journal on Computing 26, no. 5 (1997): 1484–1509. https://doi.org/10.1137/S0097539795293172.

[2]. National Institute of Standards and Technology. FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard. Gaithersburg, MD: NIST, August 13, 2024. https://csrc.nist.gov/pubs/fips/203/final.

[3]. National Institute of Standards and Technology. FIPS 204: Module-Lattice-Based Digital Signature Standard. Gaithersburg, MD: NIST, August 2024. https://csrc.nist.gov/pubs/fips/204/final.

[4]. National Institute of Standards and Technology. FIPS 205: Stateless Hash-Based Digital Signature Standard. Gaithersburg, MD: NIST, August 2024. https://csrc.nist.gov/pubs/fips/205/final.

[5]. National Institute of Standards and Technology. NIST Special Publication 800-227: Recommendations for Key-Encapsulation Mechanisms. Gaithersburg, MD: NIST, September 18, 2025. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-227.pdf.

[6]. Perera, M., M. Mackay, M. Hashem Eiza, A. Raschellà, N. Shone, and M. K. Maheshwari. "Towards Quantum-Safe O-RAN: Experimental Evaluation of ML-KEM-Based IPsec on the E2 Interface." arXiv preprint, arXiv:2601.20378, 2026. https://doi.org/10.48550/arXiv.2601.20378.

[7]. Gupta, N., H. Alimohammadi, M. Shojafar, D. Mi, and M. N. M. Bhutta. "Solving the Post-Quantum Control Plane Bottleneck: Energy-Aware Cryptographic Scheduling in Open RAN." arXiv preprint, arXiv:2602.11820, 2026. https://doi.org/10.48550/arXiv.2602.11820.

[8]. Otero-García, P., et al. "Introducing Post-Quantum Algorithms in Open RAN Interfaces." arXiv preprint, arXiv:2501.10060, 2025. https://doi.org/10.48550/arXiv.2501.10060.3GPP TSG-SA3. Liaison Statement, Meeting 99-e, S3-201xyz. May 2020.

[9]. Bhatnagar, A. "6G Standardization Intelligence Brief (2025–2030)." Medium, November 2025."3GPP Freezes Release 19 Specifications and Advances 6G Standardisation Efforts." The Critical Communications Review, March 2026.

[10]. Ericsson. "Migration to Quantum-Resistant Cryptography." Ericsson Technology Review, 2026. https://www.ericsson.com/en/reports-and-papers/ericsson-technology-review/articles/migrating-telecom-to-quantum-resistant-cryptography.

[11]. wolfSSL. "Benchmarks for Kyber Level 1 PQM4 Integration on STM32 ARM Cortex-M4." October 2024. https://www.wolfssl.com/benchmarks-kyber-level-1-pqm4-integration-stm32-arm-cortex-m4/.

[12]. Kannwischer, M. J., R. Krausz, J. Petri, and S. Yang, et al. "pqm4: Post-Quantum Crypto Library for the ARM Cortex-M4." GitHub repository, benchmarks.md. https://github.com/mupq/pqm4.

[13]. "Benchmarking NIST-Standardised ML-KEM and ML-DSA on ARM Cortex-M0+: Performance, Memory, and Energy on the RP2040." arXiv preprint, arXiv:2603.19340, 2026. https://doi.org/10.48550/arXiv.2603.19340.
[14]. "Post Quantum Migration of Tor." arXiv preprint, arXiv:2503.10238, 2025. https://doi.org/10.48550/arXiv.2503.10238.
[15]. Bos, J. W., M. Gourjon, J. Renes, T. Schneider, and C. van Vredendaal. "Masking Kyber: First- and Higher-Order Implementations." IACR Transactions on Cryptographic Hardware and Embedded Systems 2021, no. 4 (2021): 173–214.
[16]. "China's 2000-km Quantum Link Is Almost Complete." IEEE Spectrum, updated 2021.
[17]. "Implementation of Carrier-Grade Quantum Communication Networks over 10000 km." npj Quantum Information 11, article 89 (2025). https://doi.org/10.1038/s41534-025-01089-8.
[18]. "Progress on Integrating Quantum Communications in Optical Systems Testbeds." arXiv preprint, arXiv:2311.16772, 2023. https://doi.org/10.48550/arXiv.2311.16772.
[19]. "Quantum Key Distribution Spans 303 km Over Live Swedish Fiber: EuroQCI Blueprint Confirmed." Tech Times, June 9, 2026.
[20]. National Security Agency. The Commercial National Security Algorithm Suite 2.0 and Quantum Computing FAQ. U/OO/194427-22, PP-24-4014, Version 2.1. Fort Meade, MD: NSA, December 2024.
[21]. "CNSA 2.0: Complete Guide to NSA's PQC Requirements." postquantum.com, May 14, 2026.
[22]. QuSecure. "CNSA 2.0 Explained: PQC Requirements, Timelines, and Federal Impact." April 2026.
JPMorgan Chase, Toshiba, and Ciena. "JPMorgan Chase, Toshiba and Ciena Build the First Quantum Key Distribution Network Used to Secure Mission-Critical Blockchain Application." Press release, February 2022. https://www.jpmorganchase.com/newsroom/press-releases/2022/jpmc-toshiba-ciena-build-first-quantum-key-distribution-network.
[23]. JPMorgan Chase. "JPMorgan Chase Establishes Quantum-Secured Crypto-Agile Network." May 2024. https://www.jpmorgan.com/technology/news/firm-establishes-quantum-secured-crypto-agile-network.
[24]. European Commission. "European Quantum Communication Infrastructure – EuroQCI." Shaping Europe's Digital Future, 2026. https://digital-strategy.ec.europa.eu/en/policies/european-quantum-communication-infrastructure-euroqci.
[25]. European Commission. "Cable, 5G and Quantum Communications Set for €389 Million Investment from the CEF Digital Programme." November 2025.
[26]. "EU Revives Quantum Communication Investment with $97M to Secure Europe's Infrastructure." The Quantum Insider, October 2024.
[27]. European Commission. Testing and Evaluation Infrastructure for European Quantum Communication Infrastructure (EuroQCI). Call for tenders, 2023.