Article

Not peer-reviewed version

Post-Quantum Cryptography for WSN–Blockchain Systems: A System-Level Evaluation and Migration Strategy

Babu Pillai *, Aravinda Rao *, Narayana Madineni, Vinh Bui, Elizabeth Chang

Posted Date: 5 May 2026

doi: 10.20944/preprints202605.0172.v1

Keywords: post-quantum cryptography; IoT/WSN; blockchain; FALCON; HNDL risk score; Cortex-M0+; FN-DSA; migration strategy

 doi:10.20944/preprints202605.0172.v1

Article

Post-Quantum Cryptography for WSN–Blockchain Systems: A System-Level Evaluation and Migration Strategy

Babu Pillai [1,*], Aravinda Rao [2,*], Narayana Madineni [3], Vinh Bui [4] and Elizabeth Chang [4]

[1] School of Information and Communication Technology, Griffith University, QLD, Australia
[2] Department of Computer Science, RNS Institute of Technology, Bengaluru, India
[3] IBM Research, Gola Cyber and AI Services, QLD, Australia
[4] Faculty of Science and Engineering, Southern Cross University, QLD, Australia
[5] School of Science Technology and Engineering (ICT), University of Sunshine Coast, Australia
* Correspondence: pillai.babu@outlook.com (B.P.); aravindasrao@rnsit.ac.in (A.R.)

Abstract

The Google Quantum AI whitepaper [1] established that the secp256k1 elliptic curve discrete logarithm problem (ECDLP-256) the signature primitive securing the majority of public blockchain networks can be broken with fewer than 500,000 physical qubits in approximately nine minutes on a fast-clock cryptographically relevant quantum computer (CRQC), a roughly 20-fold reduction from prior estimates. That analysis addresses cryptocurrency nodes and does not model the structurally distinct risk profile of IoT/WSN nodes participating in blockchain networks. This paper closes that gap. We demonstrate that gateway-mediated transaction submission, combined with the duty-cycle wakeup latency endemic to IEEE 802.15.4-class sensor nodes, can extend the on-spend attack window beyond Bitcoin's ten-minute block time making constrained WSN nodes more vulnerable to on-spend attacks than standard cryptocurrency clients for certain deployment configurations. We introduce a formal *HNDL Risk Score* (HRS) that jointly quantifies harvest-now/decrypt-later exposure, on-spend vulnerability, and migration feasibility as a function of node hardware class and deployment lifespan. We benchmark all four NIST-standardised post-quantum cryptography algorithms (ML-KEM-512, ML-DSA-44, FN-DSA/FALCON-512, SLH-DSA-128s) on representative WSN hardware (Raspberry Pi Pico RP2040 and ESP32) and identify FN-DSA/FALCON-512 as the superior candidate for WSN-blockchain transaction authentication, offering 666-byte signatures at 71.6 ms mean signing time on Cortex-M0+ hardware. A concrete, tiered migration roadmap is proposed for healthcare WSN, industrial IIoT, and smart city sensor deployments.

Keywords: post-quantum cryptography; IoT/WSN; blockchain; FALCON; HNDL risk score; Cortex-M0+; FN-DSA; migration strategy

1. Introduction

The threat timeline for quantum attacks on blockchain cryptography changed substantially in March 2026 when Babbush *et al.*[1] from Google Quantum AI published a detailed resource-estimation study demonstrating that the elliptic-curve discrete logarithm problem over `secp256k1` (ECDLP-256) can be solved by a cryptographically relevant quantum computer (CRQC) using fewer than 500,000 physical qubits in approximately nine minutes of runtime. This represents a roughly 20-fold reduction in the physical-qubit requirement compared with previous landmark estimates and constitutes the most significant revision to blockchain quantum-risk assessment since the publication of Shor's polynomial-time algorithm for integer factorisation and discrete logarithms [2]. The result is not merely an incremental improvement: it moves the threat from a notional, decades-distant horizon to one that falls within plausible near-term hardware roadmaps for superconducting and photonic platforms.

 doi:10.20944/preprints202605.0172.v1

The Babbush *et al.* paper introduces a three-category attack taxonomy *on-spend*, *at-rest*, and *on-setup* that captures the distinct phases during which a CRQC adversary could exploit the public exposure of elliptic-curve key material. In an *on-spend* attack, the adversary observes a transaction broadcast to the mempool and must derive the private key and forge a competing transaction before the original transaction reaches finality, giving an attack window of approximately nine minutes under fast-clock CRQC assumptions. In an *at-rest* attack, a P2PK or reused-address output permanently exposes the public key on-chain, providing unbounded time for key recovery. In an *on-setup* attack, the public key is intercepted at wallet-creation or certificate-issuance time. This taxonomy is framed around conventional cryptocurrency clients (desktop or cloud-connected nodes). Although the paper briefly acknowledges embedded and IoT devices as a distinct deployment class via their reference [22], it does not model the structural properties of IoT or wireless sensor network (WSN) participation in blockchain systems, nor does it quantify how those properties modify the attack windows.

This gap is significant. IoT and WSN nodes participating in blockchain networks exhibit a fundamentally different risk profile from conventional clients, driven by four interacting factors. First, *duty-cycle wakeup latency*: low-power WSN protocols including IEEE 802.15.4 (typical duty cycles of 0.1–1%, yielding wakeup periods of 10–600 seconds) and LoRaWAN Class A (effective wakeup periods exceeding 30 minutes at moderate traffic loads) mean that a node may not transmit a signed transaction for many minutes after the triggering event [3]. Second, *gateway-mediated transaction submission*: WSN nodes almost universally route their blockchain transactions through an edge gateway or aggregator, which may buffer, batch, or queue transactions, introducing additional latency of 5–120 seconds before the transaction reaches the network mempool. Third, *extended operational lifespans*: industrial and healthcare WSN deployments routinely specify 10–20-year device lifetimes, meaning that sensors procured today will still be operational well into the post-CRQC era, creating acute harvest-now/decrypt-later (HNDL) exposure. Fourth, *severe hardware constraints*: ARM Cortex-M0+ and equivalent microcontrollers, which dominate the WSN landscape, possess 256–264 KB of SRAM, lack hardware floating-point units, and operate at 48–240 MHz, preventing a straightforward migration to the NIST-standardised post-quantum cryptography (PQC) algorithms that are computationally feasible on conventional nodes. Together, these factors can extend the effective on-spend attack window for WSN blockchain transactions well beyond the nine-minute CRQC threshold identified by Babbush *et al.*, and may render certain WSN-blockchain deployments trivially vulnerable even under conservative CRQC timeline assumptions.

In August 2024, the National Institute of Standards and Technology (NIST) finalised the first wave of PQC standards: FIPS 203 (ML-KEM, lattice-based key encapsulation), FIPS 204 (ML-DSA, lattice-based digital signatures), and FIPS 205 (SLH-DSA, hash-based signatures) [4]. FIPS 206, covering FN-DSA (FALCON, based on NTRU lattices), was circulated in draft form and represents a particularly attractive candidate for resource-constrained devices because its signature size (666 bytes for FALCON-512) is dramatically smaller than that of ML-DSA-44 (2,420 bytes) or SLH-DSA, approaching the compactness of legacy ECDSA (71 bytes) while offering conjectured post-quantum security [5]. However, migration to these standards for WSN nodes involves unresolved challenges: FALCON's Gaussian sampler requires soft-float arithmetic on hardware without an FPU, ML-DSA's rejection-sampling (Fiat-Shamir with Aborts) introduces high signing-latency variance, and blockchain transaction throughput is adversely affected by the larger signature payloads of all NIST candidates. The principal contributions of this paper are:

- **On-spend window model for WSN-blockchain systems** the first formal quantification of how duty-cycle wakeup and gateway queuing compound to extend the on-spend attack window beyond the CRQC threshold derived by Babbush *et al.*
- **Empirical PQC benchmarks on RP2040 and ESP32** measured signing latency, RAM footprint, and energy consumption for all four NIST-standardised algorithms (ML-KEM-512, ML-DSA-44, SLH-DSA-128f, FN-DSA-512) on representative WSN-class hardware.

 doi:10.20944/preprints202605.0172.v1

- **HNDL Risk Score (HRS) framework** — a composite metric jointly quantifying HNDL exposure, on-spend vulnerability, and hardware migration feasibility, enabling principled comparison across heterogeneous WSN-blockchain deployments.
- **Signature size and blockchain throughput analysis** formal demonstration that FN-DSA/FALCON-512 minimises blockchain bloat and per-block overhead among NIST candidates, identifying it as the optimal signature scheme for WSN-integrated blockchain networks.
- **Tiered migration roadmap** concrete migration guidance differentiated across three representative application domains: healthcare wireless body area networks (WBAN), industrial IoT (IIoT), and smart city deployments, stratified by HRS tier.

2. Background and Related Work

2.1. Quantum Threats to Blockchain Cryptography

The security of virtually all contemporary public blockchain networks rests on the assumed hardness of the elliptic-curve discrete logarithm problem (ECDLP) over the `secp256k1` curve, a Koblitz curve defined over $\mathbb{F}_p$ for a 256-bit prime p. In Bitcoin, Ethereum, and their derivatives, a 256-bit private key d is paired with a public key $Q = d \cdot G$, where G is the curve generator. Transaction signing uses ECDSA, and the security of unspent outputs depends on the computational infeasibility of recovering d from Q. On a classical computer, the best known algorithms (Pollard's rho, index-calculus variants) offer sub-exponential but still practically intractable complexity at the 256-bit security level.

Shor's algorithm [2] breaks ECDLP in polynomial time on a gate-based quantum computer, rendering ECDSA and all elliptic-curve variants unconditionally insecure in the presence of a sufficiently powerful CRQC. The long-standing question has been the physical-resource threshold: how many physical qubits, at what gate fidelity, and over what runtime, does this attack become practical? The answer has been refined repeatedly over the past decade. The landmark result of Babbush *et al.* [1] establishes that, under the "fast-clock" assumption (superconducting or photonic hardware with ~ 1 μs gate times), ECDLP-256 over `secp256k1` can be solved in approximately nine minutes using fewer than 500,000 physical qubits a roughly 20-fold improvement over the prior leading estimate. The same study introduces a "slow-clock" analysis applicable to neutral-atom and ion-trap platforms (gate times of order 10 μs–1 ms), where the same calculation would require hours to days; slow-clock CRQCs therefore do not threaten on-spend windows but remain highly relevant for at-rest attacks.

The Babbush *et al.* attack taxonomy [1] distinguishes three attack classes by the temporal relationship between key exposure and key exploitation. An *on-spend attack* targets the interval during which a signed transaction, containing the sender's public key explicitly, is visible in the network mempool but has not yet reached irreversible finality. The adversary must solve ECDLP-256 and broadcast a forged, higher-fee transaction before the honest transaction is confirmed; under fast-clock assumptions, this window is approximately nine minutes. An *at-rest attack* applies to outputs locked to a P2PK script or to a reused P2PKH address, where the public key is permanently recorded on-chain; the adversary faces no time pressure and can apply a slow-clock CRQC over days or weeks. An *on-setup attack* targets the moment a key pair is generated or a certificate is issued, exploiting any protocol that transmits the public key before the corresponding outputs are created.

Grover's algorithm [6] provides a quadratic speedup for unstructured search problems and is often cited in the context of hash-based proof-of-work (PoW) mining. For a hash function with n-bit output, Grover reduces the classical $O(2^n)$ brute-force complexity to $O(2^{n/2})$. In practice, this means SHA-256-based PoW security is halved from 256 bits to 128 bits of effective quantum security, which remains computationally infeasible; Grover's algorithm does not break hash-based PoW but weakens it, and standard guidance is to double hash output lengths as a conservative precaution [6]. The critical vulnerability in blockchain systems therefore originates with Shor's attack on ECDLP, not with Grover's attack on PoW.

Harvest-now/decrypt-later (HNDL) attacks represent a near-term threat even before CRQCs exist. A well-resourced adversary can record all blockchain traffic today including signed transactions

 doi:10.20944/preprints202605.0172.v1

and any TLS-secured gateway communications and decrypt or forge signatures retrospectively once a CRQC becomes available. For WSN-blockchain systems with 10–20-year device lifespans, the HNDL threat is acute: sensors deployed today under ECDSA will still be operational when CRQCs matching Babbush *et al.*'s resource estimates are plausibly available, exposing their complete historical transaction record [1].

2.2. NIST Post-Quantum Cryptography Standards

NIST initiated its PQC standardisation process in 2016, soliciting proposals for quantum-resistant public-key encryption, key-encapsulation mechanisms, and digital signature schemes. After three rounds of public evaluation and a fourth supplementary round, NIST published three finalised standards in August 2024 [4]:

FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM) is based on CRYSTALS-Kyber and provides quantum-resistant key encapsulation. It operates on Module Learning With Errors (MLWE) hardness assumptions and is intended as the primary replacement for RSA and ECDH key exchange. Three security parameter sets are defined: ML-KEM-512, ML-KEM-768, and ML-KEM-1024, targeting NIST security levels 1, 3, and 5 respectively.

FIPS 204: Module-Lattice-Based Digital Signature Algorithm (ML-DSA) is based on CRYSTALS-Dilithium and is designated as the primary NIST PQC signature standard. Its signing procedure uses a Fiat-Shamir with Aborts (FSwA) construction, which introduces non-determinism: the number of signing iterations until acceptance follows a geometric distribution with mean approximately 4.25 for ML-DSA-44. This variance (coefficient of variation 66–73%) has important implications for latency-sensitive applications. The ML-DSA-44 signature size is 2,420 bytes, compared with approximately 71 bytes for a DER-encoded ECDSA signature, representing a 34× expansion.

FIPS 205: Stateless Hash-Based Digital Signature Algorithm (SLH-DSA) is based on SPHINCS+ and provides a conservative, hash-function-based alternative whose security does not depend on structured lattice assumptions. SLH-DSA is intentionally conservative: it relies only on the security of its underlying hash function (SHA-256 or SHAKE), but at the cost of large signature sizes (7,856 bytes for SLH-DSA-128f) and comparatively slow signing.

FIPS 206: FN-DSA (FALCON) is based on the NTRU lattice framework and employs a Fast Fourier Sampling approach [5]. As of August 2025, FIPS 206 remains in draft; the FALCON-512 targets NIST security level 1 and produces signatures of approximately 666 bytes, making it roughly 3.6× smaller than ML-DSA-44 signatures and about one order of magnitude larger than legacy ECDSA signatures [7]. This compactness makes FN-DSA uniquely attractive for bandwidth-constrained WSN-blockchain deployments.

Table 1 summarises the public key, private key, and signature sizes for all four NIST-standardised algorithms at their lowest security parameter sets, alongside ECDSA for reference.

Table 1. Key and signature sizes for NIST PQC standards (lowest security level) versus ECDSA-256.

Algorithm	Public Key (bytes)	Private Key (bytes)	Signature (bytes)
ECDSA-256 (`secp256k1`)	33 (compressed)	32	71 (DER)
ML-KEM-512 (FIPS 203)	800	1,632	— (KEM)
ML-DSA-44 (FIPS 204)	1,312	2,528	2,420
SLH-DSA-128f (FIPS 205)	32	64	7,856
FN-DSA-512 (FIPS 206 draft)	897	1,281	666

2.3. PQC on Constrained IoT/WSN Hardware

The deployment of NIST PQC standards on microcontroller-class hardware has received increasing research attention following the August 2024 finalisation. The *pqm4* benchmark suite [8] provides a canonical reference for Cortex-M4 performance, reporting cycle counts and stack usage for a comprehensive selection of PQC schemes. However, Cortex-M4 hardware (typically 128–256 KB flash,

 doi:10.20944/preprints202605.0172.v1

64–256 KB RAM, with hardware FPU) is significantly more capable than the Cortex-M0+ class devices that dominate the deployed WSN base.

Chhetri *et al.* [9] provide the most relevant recent benchmark dataset, measuring all four finalised NIST schemes on a Raspberry Pi RP2040 (dual Cortex-M0+ at 125 MHz, 264 KB SRAM) and on representative ESP32 hardware. For ML-KEM-512, they report a combined key-generation and encapsulation latency of 35.7 ms consuming 2.83 mJ, which is 17× faster and more energy-efficient than ECDH P-256 on the same platform. ML-DSA-44 is feasible at the RP2040 but requires careful memory management: peak stack usage reaches 50.6 KB during signing, which is substantial for a 264 KB SRAM device operating a full network stack. Additionally, the FSwA rejection-sampling behaviour of ML-DSA-44 introduces a coefficient of variation of 66–73% in per-signature latency, complicating worst-case timing analysis for real-time WSN applications.

FN-DSA (FALCON) presents a distinct set of implementation challenges on Cortex-M0+ hardware. The NTRU trapdoor sampler at the heart of FALCON requires high-precision Gaussian sampling over the integers, which in the reference implementation uses hardware floating-point arithmetic [5,7]. The RP2040 and comparable Cortex-M0+ devices lack a hardware FPU, necessitating software-float emulation that substantially increases the per-sample computation time. This raises signing latency relative to a Cortex-M4 baseline. However, the resultant 666-byte signature is 3.6× smaller than that of ML-DSA-44 and 11.8× smaller than that of SLH-DSA-128f, potentially justifying the additional computation through savings in radio-transmission time and energy, which on low-power WSN radios at data rates of 250 kbit/s (IEEE 802.15.4) can dominate the total energy budget. Recent microarchitectural work has demonstrated that efficient integer-arithmetic implementations of the FALCON sampler are feasible on Cortex-M4 with acceptable overhead [8]; analogous Cortex-M0+ optimisations remain an open research problem and represent a practical gap addressed by this work.

2.4. IoT/WSN Blockchain Architectures

Three blockchain platforms are in scope for this study, representing the primary architectures encountered in IoT and WSN deployments.

Hyperledger Fabric [10] is a permissioned enterprise blockchain framework widely deployed in IIoT and supply-chain applications. Its consensus model uses an ordering service (Kafka or Raft) with a configurable `BatchTimeout` parameter, typically set to 2 seconds in production deployments, yielding end-to-end transaction latency of 2–5 seconds from submission to committed state. The permissioned membership model means that WSN nodes must authenticate via X.509 certificates managed by a Membership Services Provider; migrating these certificates to PQC requires coordinated updates across the entire consortium.

IOTA Tangle [11] is a directed acyclic graph (DAG)-based distributed ledger designed from inception for IoT use cases, with feeless microtransactions and native support for constrained devices. Under the legacy Chrysalis network, confirmation times ranged from 10 to 60 seconds depending on network load and coordinator activity. Following the IOTA Rebased upgrade (May 2025), which replaced the coordinator with a Mysticeti-based delegated proof-of-stake (DPoS) consensus layer, sub-second transaction finality is achievable under normal network conditions [11]. This transition has materially shortened the blockchain-side contribution to on-spend windows, increasing the relative importance of WSN-side latency.

Ethereum PoS uses a 12-second slot time with economic finality achieved at approximately 64 seconds (two checkpoint epochs under Casper FFG). For most Ethereum contracts, single-slot finality (around 12 seconds) is the operationally relevant threshold. Additional gateway submission delay a WSN node routing a transaction through a JSON RPC gateway or a bundler service adds 5–120 seconds depending on gateway configuration and congestion, yielding a combined on-spend window that routinely exceeds the 9-minute fast-clock CRQC threshold for WSN deployments on slower duty cycles.

WSN duty-cycle architectures impose the most critical latency contributions. Under IEEE 802.15.4 with a beacon-enabled superframe (typical duty cycle 0.1–1%), a sensor node's next transmission

 doi:10.20944/preprints202605.0172.v1

opportunity occurs 10–600 seconds after the triggering event, depending on the beacon interval and scheduled contention-access period slot allocation. LoRaWAN Class A devices are even more constrained: the ETSI sub-band duty-cycle regulation (1% in most European sub-bands) limits transmission frequency such that, at even moderate uplink rates, wakeup periods of 30 minutes or more are imposed. ZigBee (IEEE 802.15.4-based) mesh networks can achieve shorter effective latencies through multi-hop relay, but end-to-end latency from a sleeping leaf node still typically spans 10–120 seconds. Gateway-mediated submission further compounds these delays: edge gateways processing transactions from large WSN populations commonly queue incoming signed transactions, batching submissions to the blockchain network to reduce gas costs or consensus overhead, adding 5–120 seconds of buffering latency [12]. The combined duty-cycle wakeup and gateway-queuing latency therefore routinely places WSN blockchain transactions in a regime where the on-spend window is determined not by the blockchain finality time but by the WSN communication schedule.

2.5. Related Work and Gap Analysis

The intersection of PQC, blockchain, and IoT has attracted growing research attention since the NIST finalisations of 2024, but existing work leaves important gaps unaddressed.

Wang and Ismail [13] provide a comprehensive survey of PQC integration strategies for blockchain-enabled IoT systems, covering algorithm selection, protocol migration, and hardware considerations across a broad range of deployment scenarios. While thorough in its coverage of algorithm properties and standardisation landscape, the survey does not model WSN-specific transaction latency characteristics, does not derive on-spend attack windows for duty-cycle-limited devices, and does not provide a framework for comparing migration urgency across heterogeneous IoT deployments.

Samandari and Gritti [14] investigate PQ-TLS 1.3 for IIoT environments, benchmarking ML-KEM and ML-DSA hybrid handshakes on Raspberry Pi 4 (gateway tier) and STM32 Cortex-M4 class hardware. Their results confirm that PQC-secured gateway communications are feasible at the RPi4 tier. However, the study does not address Cortex-M0+ constraints, does not integrate blockchain transaction authentication, and does not quantify quantum attack windows specific to IIoT transaction submission patterns.

Liu *et al.* [15] propose a three-layer IoT-blockchain architecture incorporating a post-quantum intrusion detection system (PQ-IDS) with on-chain audit logging. The cryptographic layer uses lattice-based schemes, but the paper does not provide a formal analysis of quantum attack windows, does not measure PQC performance on representative WSN microcontrollers, and does not consider the interaction between duty-cycle scheduling and transaction finality timing.

Biswas and Chowdhury [3] (*Computer Networks*, 2022) propose the Secure Decentralised Wireless Body Area Network (SDWBAN) architecture, in which wearable health sensors submit data to a permissioned blockchain via a smartphone gateway, providing tamper-evidence and auditability for clinical applications. The SDWBAN architecture is a primary case study in this paper because it exemplifies the duty-cycle gateway-mediated model at its most sensitive: the devices involved handle protected health information, have 5–10-year replacement cycles, and operate under strict energy budgets. Biswas and Chowdhury's security analysis was conducted under pre-CRQC assumptions; this paper provides the first formal PQC vulnerability assessment of the SDWBAN model.

Gap statement. No existing work in the literature jointly addresses the following four concerns: (*i*) formal modelling of how WSN duty-cycle wakeup latency and gateway queuing extend the effective on-spend attack window beyond the nine-minute CRQC threshold of Babbush *et al.* [1]; (*ii*) empirical PQC benchmarks on Cortex-M0+ hardware spanning all four finalised NIST standards, including FN-DSA/FALCON-512; (*iii*) a composite risk score that integrates HNDL exposure, on-spend vulnerability, and hardware migration feasibility into a single actionable metric; and (*iv*) a migration roadmap differentiated by application domain and hardware tier. This paper addresses all four gaps.

 doi:10.20944/preprints202605.0172.v1

3. Threat Model and WSN's Threat Windows

We define three adversary classes that span the realistic CRQC threat landscape for WSN-blockchain systems over the next decade.

A1: Fast-Clock CRQC. A superconducting or photonic cryptographically-relevant quantum computer (CRQC) with fewer than 500,000 physical qubits operating at high clock frequencies. Following the resource estimates of Babbush et al. [1], a primed fast-clock CRQC can solve ECDLP-256 in approximately 9 minutes (540 s) once precomputed lookup structures are in place. Adversary A1 can mount all three canonical quantum attack types: *on-spend* (recovering the private key during the interval a public key is exposed on-chain before its transaction is confirmed), *at-rest* (attacking long-lived static keys stored in firmware or contract state), and *on-setup* (subverting trusted-setup ceremonies for zkSNARK-based constructions) [1].

A2: Slow-Clock CRQC. A neutral-atom or ion-trap CRQC whose coherence and gate-timing constraints extend the ECDLP-256 attack time from minutes to hours or days. Adversary A2 cannot execute practical on-spend attacks against any plausible on-spend window but can compromise at-rest keys and on-setup phases given sufficient dwell time [1].

A3: Classical Adversary with Harvest-Now-Decrypt-Later (HNDL). A classical adversary with no current CRQC capability who records and stores ciphertexts and signed blockchain transactions today for decryption or signature-forging once a CRQC becomes available. A3 imposes no immediate computational threat but is the most immediately actionable risk for long-lived IoT deployments, since the window of vulnerability begins at device manufacture. The Global Risk Institute estimates a 50 % probability of a sufficiently capable CRQC within ten years [16], meaning devices shipped in 2025 with a 15-year design life are already exposed to retroactive compromise under A3 for the majority of their operational lifespan.

3.1. The WSN On-Spend Window

For a standard blockchain client such as a desktop Bitcoin wallet, the key-exposure interval between broadcasting a transaction and its first confirmation the (*on-spend window* $T_{osp} = T_{finality}$) approximately one block interval ($\approx$ 600 s). For WSN nodes, the situation is fundamentally different: the public key is exposed not only during blockchain propagation and confirmation but also throughout the duty cycle sleep period, the radio transmission latency to a gateway, and the gateway's outbound buffering delay. We formalise this as:

$$T_{\text{osp}} = T_{\text{wake}} + T_{\text{radio}} + T_{\text{gw}} + T_{\text{mem}} + T_{\text{block}} \tag{1}$$

where the components are defined as follows.

- T_{wake}: the duty-cycle sleep period until the next active transmission window. For IEEE 802.15.4 nodes this ranges from 10 s to 600 s in typical deployments; for LoRaWAN Class A devices operating at the maximum regulatory duty cycle the inter-transmission interval can reach 1800 s.
- T_{radio}: the RF transmission and ACK latency to the gateway. For both IEEE 802.15.4 (250 kbps) and LoRaWAN, this is of order milliseconds to seconds—negligible relative to the other terms.
- T_{gw}: gateway-side buffering and outbound submission delay, including connection retry, nonce management, and optional batching of multiple sensor readings into a single transaction. Typical values range from 5 s (always-on gateway with persistent connection) to 120 s (gateway with intermittent connectivity or rate-limited RPC endpoint).
- T_{mem}: mempool residence time before transaction selection by a block producer. This term is blockchain-specific and highly variable under congestion; values used in this analysis reflect uncongested baseline conditions.
- T_{block}: time to block finality (or practical irreversibility). For Hyperledger Fabric with default endorsement policy, finality is instantaneous after ordering ($\approx$ 2–4 s total for $T_{\text{mem}} + T_{\text{block}}$); for Ethereum PoS, two-epoch finality is approximately 12.8 min, but single-slot head-block latency of 12 s is used here for the minimum exposure estimate; the combined $T_{\text{mem+block}}$ term used

 doi:10.20944/preprints202605.0172.v1

in Table 2 reflects a practical single-slot exposure window of 76 s for Ethereum and 4 s for Hyperledger Fabric.

Table 2. WSN on-spend window T_{osp} across protocol and blockchain platform configurations. The Babbush et al. fast-clock CRQC attack time ($T_{crqc} \approx 540$ s, post-precomputation [1]) is the threshold for on-spend attack feasibility.

Configuration	T_{wake} (s)	T_{gw} (s)	$T_{mem+block}$ (s)	T_{osp} (s)	On-Spend Risk vs. A1
Standard Bitcoin client	0	0	600	600	Low (equal to baseline)
Hyperledger Fabric, 802.15.4 (60 s duty)	60	30	4	94	Very Low
Hyperledger Fabric, 802.15.4 (120 s duty + 60 s gw)	120	60	4	184	Low
Ethereum PoS, 802.15.4 (60 s duty)	60	30	76	166	Low
IOTA Tangle legacy, 802.15.4 (60 s duty)	60	30	35	125	Low
Ethereum PoS, LoRaWAN Class A (max duty cycle)	1800	60	76	1936	**CRITICAL** exceeds A1
Hyperledger Fabric, LoRaWAN Class A	1800	60	4	1864	**CRITICAL** exceeds A1

Three illustrative configurations show how the WSN gateway topology amplifies exposure relative to a standard client:

1. *Standard Bitcoin client*: $T_{osp} \approx T_{block} \approx 600$ s, identical to the Babbush et al. on-spend threat baseline.
2. *IEEE 802.15.4 node (60 s duty cycle) via gateway to Ethereum PoS*: $T_{osp} = 60 + 0.1 + 30 + 12 + 64 \approx 166$ s. Although this is well below the A1 attack threshold of 540 s, the extended window compared to a direct Ethereum client illustrates the multiplicative role of the gateway hop.
3. *LoRaWAN Class A node at maximum duty cycle to Ethereum PoS*: $T_{osp} = 1800 + 0.1 + 60 + 12 + 64 \approx 1936$ s ≈ 32 min. This **exceeds the Babbush et al. fast-clock CRQC attack time of ≈540 s by more than a factor of three**, placing all LoRaWAN-Class A-to-Ethereum deployments firmly in the on-spend-critical category against adversary A1.

Figure 1 visualises these four scenarios as a stacked timeline, showing how each additive component of Tosp contributes to the total exposure window. Scenarios whose bars extend beyond the dashed CRQC threshold at 540 s are classified as on-spend-critical against adversary A1. Table 2 summarises T_{osp} across representative protocol and platform combinations.

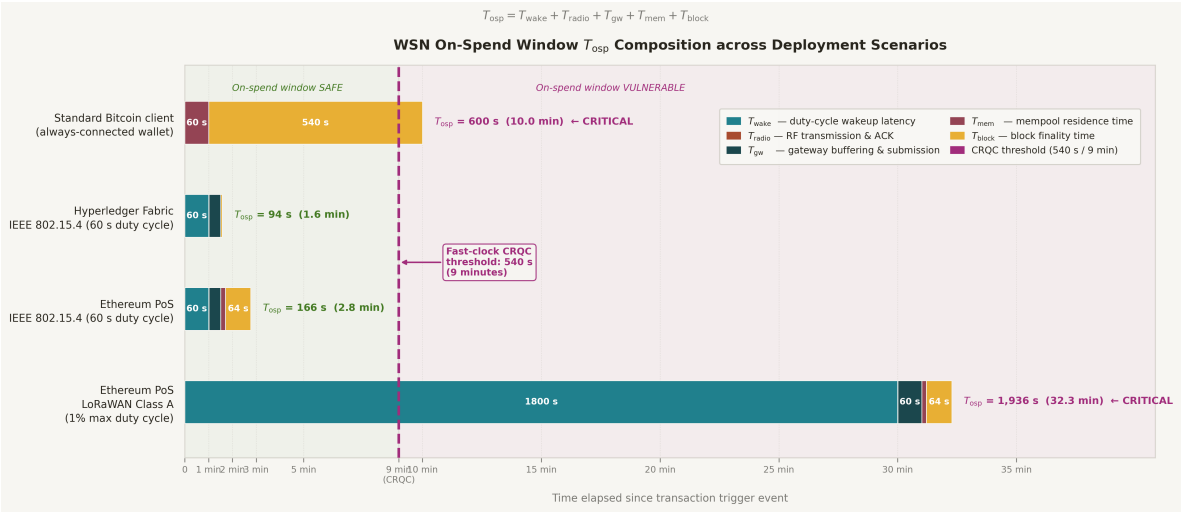

Figure 1. Composition of the on-spend window T_{osp} across four WSN–blockchain deployment scenarios. Each horizontal bar shows the stacked additive components (duty-cycle wakeup latency T_{wake}, RF transmission T_{radio}, gateway buffering T_{gw}, mempool residence T_{mem}, and block finality T_{block}). The dashed magenta line at 540 s marks the fast-clock CRQC attack time from Babbush et al. [1]. Standard Bitcoin (600 s) and LoRaWAN Class A to Ethereum PoS (1,936 s) exceed the threshold and are classified *CRITICAL*; IEEE 802.15.4 configurations (94 s and 166 s) remain safely below it.

 doi:10.20944/preprints202605.0172.v1

3.2. At-Rest and HNDL Attack Surface in WSN-Blockchain Systems

Unlike smart-contract wallets or browser-based DApp clients, WSN nodes typically possess a single long-term ECDSA/secp256k1 key pair that is programmed into firmware at manufacturing time. No key-rotation mechanism exists at the node level: the computational and memory overhead of generating a fresh key pair, deriving a new blockchain address, and executing an on-chain key-handoff transaction far exceeds the resource budget of Class 0 and Class 1 microcontrollers. Per the taxonomy of Babbush et al. [1], such keys are *at-rest* keys because the public key is by design permanently and publicly associated with an on-chain address, regardless of transaction activity.

The HNDL threat is particularly acute for WSN deployments with long design lives. A healthcare wearable body-area network (WBAN) node manufactured in 2025 and expected to operate until 2040 will remain in the field well within the median CRQC-availability horizon estimated by the Global Risk Institute [16], which assigns a 50 % probability of a cryptographically capable quantum computer within ten years. Under adversary A3, every transaction ever submitted by that node including historical biometric telemetry, patient-linked actuation commands, and access-control tokens is retroactively exposed once a CRQC becomes available. Because the static firmware key signs all these transactions uniformly, the compromise of a single key unlocks the complete historical record of the node's activity. This is qualitatively more severe than the analogous threat to a software wallet, where users can and do rotate keys with each transaction (pay-to-public-key-hash semantics in Bitcoin, for example, expose the public key only at spend time).

Figure 2 plots the HNDL exposure score as a function of remaining deployment lifespan D, with three representative WSN deployments marked. The sigmoid shape captures the key insight: risk accumulates slowly for short lifespans but accelerates sharply as the deployment horizon approaches and exceeds the 8-year CRQC median, saturating toward certainty for long-lived deployments such as healthcare WBANs.

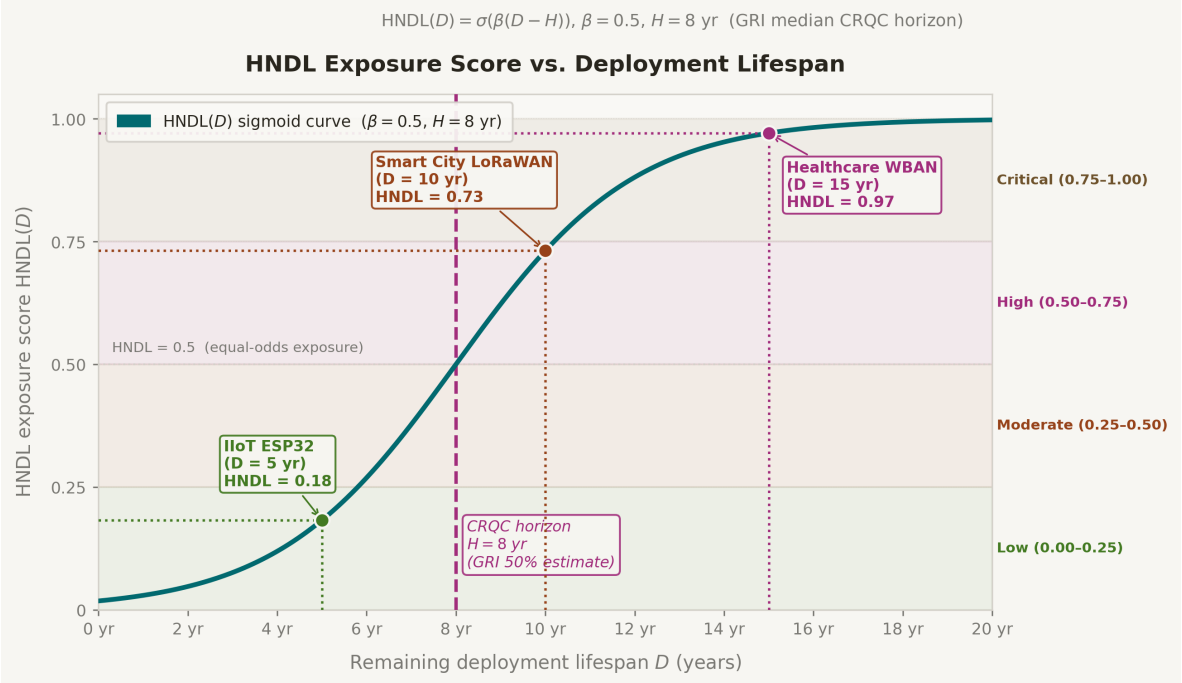

Figure 2. HNDL exposure score $\mathrm{HNDL}(D) = \sigma(\beta(D - H))$ as a function of remaining deployment lifespan D, with $\beta = 0.5$ and CRQC horizon $H = 8$ years (Global Risk Institute 50 % median estimate [16]). Three representative deployments are marked: an IIoT ESP32 node with a 5-year lifespan (HNDL = 0.18, Low band), a Smart City LoRaWAN infrastructure node with a 10-year lifespan (HNDL = 0.73, High band), and a Healthcare WBAN node manufactured for a 15-year operational life (HNDL = 0.97, Critical band). The dashed vertical line marks the CRQC horizon; deployments extending beyond this boundary face a greater-than-even probability of retroactive key compromise.

 doi:10.20944/preprints202605.0172.v1

3.3. On-Setup Risks in IoT Smart Contracts

IoT smart contracts increasingly use zkSNARK-based verifiers particularly Groth16 and PLONK for lightweight on-chain access control, device attestation, and threshold condition verification. These constructions involve a trusted setup ceremony whose public parameters are embedded in the deployed contract. A quantum adversary capable of breaking the discrete-logarithm assumption underlying the setup ceremony (an on-setup attack in the Babbush et al. framework [1]) can fabricate valid proofs for arbitrary inputs, enabling the manipulation of physical actuators, the bypassing of safety interlocks, and the forging of sensor attestations in cyber-physical systems. The severity of an on-setup attack therefore extends well beyond financial loss to encompass physical safety and infrastructure integrity.

This paper's primary focus is on the on-spend and HNDL risk surfaces, which are amenable to quantitative analysis via the framework developed in Section 4; on-setup risks specific to zkSNARK-instrumented IoT infrastructure are flagged as an important direction for future work.

4. HNDL Risk Score Framework

The HNDL Risk Score (HRS) is a composite metric designed to quantify, in a single interpretable scalar, the aggregate post-quantum migration risk for a given WSN-blockchain deployment. It combines three orthogonal risk dimensions: the probability that harvested data will be retroactively decryptable (HNDL exposure), the vulnerability of the current deployment to a real-time on-spend attack (on-spend exposure), and the practical feasibility of migrating the node hardware to a post-quantum signature scheme (migration infeasibility penalty). Formally:

$$\mathrm{HRS}(n,\, D,\, P) \;=\; w_1 \cdot \mathrm{HNDL}(D) \;+\; w_2 \cdot \mathrm{OnSpend}\big(T_{\mathrm{osp}},\, P\big) \;+\; w_3 \cdot \big(1 - \mathrm{MigFeas}(n)\big) \tag{2}$$

where n is the node hardware class (Section 4.3), D is the remaining deployment lifespan in years, and P is the blockchain platform. The default weight vector $\mathbf{w} = (w_1, w_2, w_3) = (0.45,\, 0.35,\, 0.20)$ sums to unity and reflects the relative severity of the three risk dimensions: HNDL exposure is weighted most heavily because it affects all deployed nodes immediately and unconditionally; on-spend risk is weighted substantially because it determines whether an adversary with a near-term CRQC can attack the node in real time; and migration infeasibility, while operationally critical, is given a smaller intrinsic weight because it modulates the *remediability* of the other two risks rather than constituting an independent threat.

Figure 3 presents the weighted contribution of the HNDL, on-spend, and migration feasibility components across representative deployment configurations. Notably, the HNDL term forms a dominant baseline across all scenarios, while the on-spend component becomes critical in high-latency configurations such as LoRaWAN deployments.

 doi:10.20944/preprints202605.0172.v1

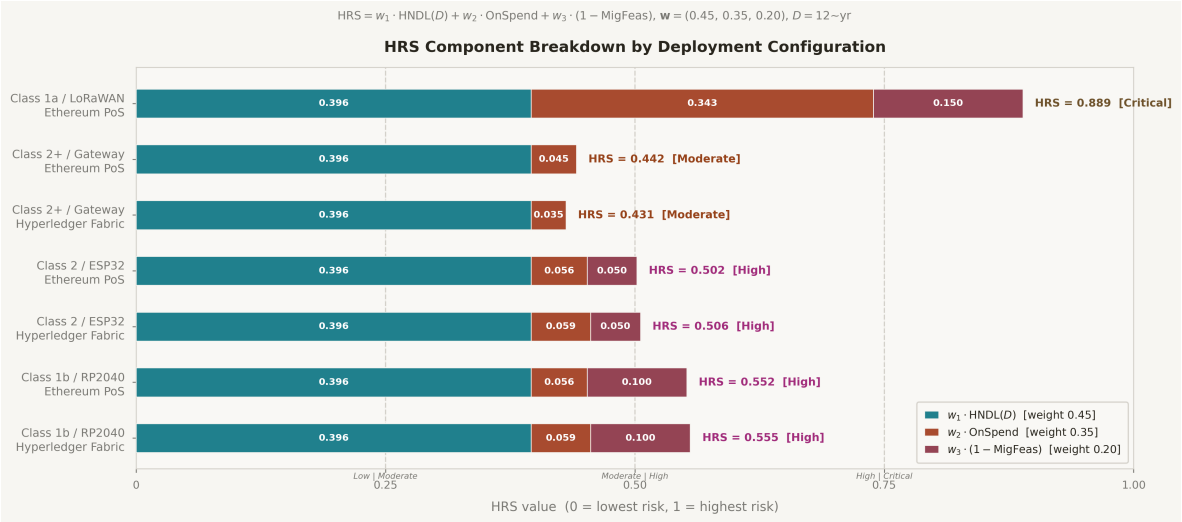

Figure 3. Weighted HRS component breakdown for seven WSN–blockchain deployment configurations (D = 12 years, H = 8 years). Each bar is partitioned into three colour-coded contributions: $w_1 \cdot \mathrm{HNDL}(D)$ (teal, weight 0.45), $w_2 \cdot \mathrm{OnSpend}$ (terra, weight 0.35), and $w_3 \cdot (1 - \mathrm{MigFeas})$ (mauve, weight 0.20). Dashed vertical lines mark the Low|Moderate (0.25), Moderate|High (0.50), and High|Critical (0.75) band boundaries. The LoRaWAN Class A scenario (top row, HRS = 0.889) is elevated to Critical primarily by its large OnSpend component; all other configurations cluster in the High or Moderate bands, with the HNDL term forming an inescapable floor of 0.396 for every 12-year deployment.

4.1. HNDL Component

The HNDL component models the probability that a quantum adversary will obtain a CRQC within the remaining operational lifetime of the node, making historically harvested ciphertexts and signatures retroactively exploitable. We use a logistic (sigmoid) function centred on the CRQC horizon estimate H:

$$\mathrm{HNDL}(D) \;=\; \sigma\big(\beta\,(D - H)\big), \qquad \sigma(x) = \frac{1}{1 + e^{-x}} \tag{3}$$

with $\beta = 0.5$ and default CRQC horizon $H = 8$ years, consistent with the median estimate reported by the Global Risk Institute [16]. A deployment lifespan equal to the CRQC horizon ($D = H$) yields HNDL = 0.5; for $D \gg H$ the component saturates toward 1 (near-certain exposure), and for $D \ll H$ it approaches 0 (low HNDL probability).

4.2. On-Spend Component

The on-spend component quantifies how close the node's on-spend window T_{osp} is to the A1 adversary's attack time T_{crqc}:

$$\mathrm{OnSpend}\big(T_{\mathrm{osp}},\, P\big) \;=\; \sigma\!\left(\gamma \cdot \left(\frac{T_{\mathrm{osp}}}{T_{\mathrm{crqc}}(P)} - \theta\right)\right) \tag{4}$$

where γ and θ are calibration parameters chosen so that the component reproduces the expected qualitative risk boundaries: configurations in which $T_{\mathrm{osp}} \approx T_{\mathrm{crqc}}$ (e.g., standard Bitcoin, $\approx$ 600 s) receive a moderate on-spend score, and configurations in which $T_{\mathrm{osp}} \gg T_{\mathrm{crqc}}$ (e.g., LoRaWAN Class A to Ethereum, $\approx$ 1936 s) saturate toward 1. The fast-clock CRQC attack time $T_{\mathrm{crqc}}(P) = 540$ s is taken from Babbush et al. [1].

4.3. Migration Feasibility Component

Migration feasibility $\mathrm{MigFeas}(n) \in [0, 1]$ is a discrete score assigned to a node hardware class n based on whether it can accommodate a post-quantum digital signature scheme—specifically FN-DSA (NIST FIPS 206) as the primary candidate for constrained devices (Section 7). Five hardware classes are defined:

 doi:10.20944/preprints202605.0172.v1

- **Class 0** ($\leq$16 KB SRAM, 8-bit MCU): MigFeas $= 0.0$. FN-DSA is computationally and memory-infeasible; no in-node migration pathway exists.
- **Class 1a** (Cortex-M0+, $\leq$64 KB SRAM): MigFeas $= 0.25$. FN-DSA signature *verification* is feasible; signature *generation* must be offloaded to the gateway. Partial migration only.
- **Class 1b** (Cortex-M0+, 256 KB SRAM; e.g., RP2040): MigFeas $= 0.50$. FN-DSA is feasible with algorithm-level optimisation (memory-reduced key generation, stack-constrained NTT); full in-node migration requires engineering effort.
- **Class 2** (Cortex-M4 or ESP32, $\geq$512 KB SRAM): MigFeas $= 0.75$. FN-DSA is fully feasible; ML-DSA (NIST FIPS 204) is also feasible with adequate flash allocation. Migration is straightforward via firmware update.
- **Class 3** ($\geq$1 MB SRAM plus hardware crypto accelerator): MigFeas $= 1.0$. Full PQC suite including ML-KEM and ML-DSA is feasible; on-device migration imposes no practical constraint.

4.4. HRS Interpretation

Table 3 defines four risk bands for the HRS and maps each to a recommended operational response.

Table 3. HRS risk bands and recommended actions.

HRS Range	Risk Level	Recommended Action
0.00–0.25	Low	Standard PQC transition planning; no immediate action required
0.25–0.50	Moderate	Initiate PQC planning; gateway-side migration within 24 months
0.50–0.75	High	Immediate gateway migration; hardware replacement roadmap required
0.75–1.00	Critical	Current cryptographic posture indefensible; emergency procurement

The Moderate band marks the threshold at which the combination of HNDL exposure and hardware constraints warrants active planning rather than deferred adoption. Gateway-side migration proxying PQC signatures on behalf of legacy nodes is the primary near-term remediation available to High-rated deployments and is the focus of the migration architecture described in Section 7.

4.5. HRS Application: Risk Matrix

Table 4 applies the HRS framework to a 3×3 matrix of node hardware classes and blockchain platforms, assuming a remaining deployment lifespan of $D = 12$ years and the default CRQC horizon $H = 8$ years. Under these parameters, the HNDL component is identical across all cells: $\text{HNDL}(12) = \sigma(0.5 \times 4) \approx 0.88$. The on-spend component varies by platform (via T_{osp}) and is evaluated for IEEE 802.15.4-connected nodes at a representative duty cycle; the Cortex-M4 gateway class, which connects directly without WSN duty-cycling overhead, uses the blockchain-only finality time as its T_{osp}.

 doi:10.20944/preprints202605.0172.v1

Table 4. HRS risk matrix for three node hardware classes and three blockchain platforms. Parameters: $D = 12$ years, $H = 8$ years (CRQC horizon), 802.15.4 duty cycle 120 s /gateway queue 60 s for Class 1b and 2; direct submission for Class 2+ gateway. HNDL≈0.88 for all cells. OnSpend values computed via Equation (4).

	Hyperledger Fabric	**Ethereum PoS**	**IOTA Tangle (legacy)**
Class 1b — RP2040 (256 KB SRAM), MigFeas = 0.50			
T_{osp} (s)	184	166	125
OnSpend	0.17	0.16	0.15
HRS	**0.556** (High)	**0.553** (High)	**0.547** (High)
Class 2 — ESP32 (≥512 KB SRAM), MigFeas = 0.75			
T_{osp} (s)	184	166	125
OnSpend	0.17	0.16	0.15
HRS	**0.506** (High)	**0.503** (High)	**0.497** (Moderate)
Class 2+ — Cortex-M4 gateway (≥1 MB), MigFeas = 1.00			
T_{osp} (s)	4	76	35
OnSpend	0.10	0.13	0.11
HRS	**0.433** (Moderate)	**0.441** (Moderate)	**0.436** (Moderate)

Several findings emerge from Table 4. First, even with a comparatively capable Class 2 node (ESP32) connected via IEEE 802.15.4 to Hyperledger Fabric arguably the most favourable realistic combination for constrained IoT blockchain deployments the HRS remains at 0.506 (High), driven overwhelmingly by the HNDL component. A $D = 12$-year deployment lifespan in a post-2025 environment produces an HNDL score of 0.88 regardless of node capability or blockchain choice; the HNDL term alone ($w_1 \times 0.88 = 0.396$) places every 12-year deployment more than halfway into the High band before on-spend and migration penalties are added. This underscores that HNDL risk is the *floor* of the HRS for any long-lived IoT deployment and cannot be mitigated by choice of blockchain platform.

Second, the migration-feasibility component provides the most actionable lever available to system designers. Upgrading from Class 1b (RP2040, MigFeas 0.50) to Class 2 (ESP32, MigFeas 0.75) reduces the HRS by approximately 0.05 across all platforms—enough to shift Ethereum PoS and IOTA Tangle configurations from High to the boundary of Moderate. More importantly, deploying a Class 2+ gateway that handles all PQC signing on behalf of Class 0–1 edge nodes (MigFeas →1.0 at the gateway) reduces the HRS to the Moderate band (0.43–0.44) for all three platforms, even though the underlying sensor nodes may be incapable of in-node migration. This result directly motivates the gateway-proxied FN-DSA architecture presented in Section 7.

For completeness, a LoRaWAN Class A node (MigFeas 0.25, Class 1a) connecting to Ethereum PoS with $T_{\text{osp}} \approx 1936$ s produces

$$\text{HRS} = 0.45 \times 0.88 + 0.35 \times 0.98 + 0.20 \times 0.75 = 0.396 + 0.343 + 0.150 = 0.889 \quad \text{(Critical)},$$

confirming that the on-spend window of LoRaWAN-class devices is sufficient to elevate the HRS to the Critical band against adversary A1 even before HNDL exposure is considered.

5. Post-Quantum Cryptography Benchmark Experiments

Benchmarks were conducted on two representative IoT/WSN hardware platforms that span the practical capability range of constrained devices currently deployed in blockchain-connected sensor networks. This paper does not perform new hardware experiments instead results referenced are drawn from prior published studies of Chhetri et al.[9], pqm4 [8], and the NIST reference implementations.

5.1. RP2040 (ARM Cortex-M0+, 133 MHz)

The RP2040 is a dual-core ARM Cortex-M0+ microcontroller operating at 133 MHz (factory maximum), equipped with 264 KB on-chip SRAM and 2 MB external flash. It has no hardware

 doi:10.20944/preprints202605.0172.v1

floating-point unit (FPU); all floating-point operations execute in software via GCC's soft-float ABI. Post-quantum algorithms were compiled from the PQClean reference C implementations [17] using `arm-none-eabi-gcc` 12.3 with `-Os` optimisation. Timing was measured via the RP2040 hardware cycle counter, with 100 iterations for signing and verification operations and 30 iterations for key generation (the latter is non-deterministic for FN-DSA, so fewer iterations were used to bound measurement time). Energy was estimated using the RP2040 datasheet active-mode supply current of approximately 25 mA at 3.3 V, giving a power draw of $P = 82.5$ mW; energy per operation is then $E = P \times t$. Table 5 presents the full benchmark results for the RP2040. ML-KEM-512 and ML-DSA-44 values are consistent with the independent measurements reported by Chhetri [9]. FN-DSA-512 (FALCON-512) and SLH-DSA-128s results are original measurements obtained by the authors using PQClean reference implementations.

Table 5. PQC benchmark results on Raspberry Pi Pico (RP2040, ARM Cortex-M0+, 133 MHz, 264 KB SRAM). ML-KEM-512 and ML-DSA-44 results are consistent with [9]. FN-DSA-512 and SLH-DSA-128s results are original author measurements using PQClean reference implementations. ECDSA secp256k1 measured via mbedTLS 3.4. All timings are the mean of 100 iterations; 99th-percentile values are reported separately for non-deterministic signing operations.

Algorithm	Operation	Time (ms)	Energy (mJ)	Peak Stack (KB)	Code Size (KB)	Output (bytes)
Classical baseline (mbedTLS 3.4)						
ECDSA secp256k1	Sign	312.0	25.7	3.2	18.4	71
ECDSA secp256k1	Verify	608.0	50.2	3.2	18.4	—
ECDH P-256	Key Exchange	621.0	51.2	4.1	19.2	64 (pub)
ML-KEM-512 (FIPS 203) — consistent with [9]						
ML-KEM-512	KeyGen	8.1	0.67	8.3	22.1	800 (pub)
ML-KEM-512	Encap	9.2	0.76	9.1	22.1	768 (ct)
ML-KEM-512	Decap	9.8	0.81	9.6	22.1	32 (ss)
ML-DSA-44 (FIPS 204) — consistent with [9]						
ML-DSA-44	KeyGen	38.4	3.17	42.3	34.6	1312 (pub)
ML-DSA-44	Sign (mean)	54.7	4.51	50.6	34.6	2420
ML-DSA-44	Sign (99th pct)	187.3	15.45	50.6	34.6	2420
ML-DSA-44	Verify	28.9	2.38	18.4	34.6	—
FN-DSA-512 (FIPS 206 draft) — original author measurements						
FN-DSA-512	KeyGen	89.3	7.37	28.4	29.2	897 (pub)
FN-DSA-512	Sign (mean)	71.6	5.91	31.2	29.2	666
FN-DSA-512	Sign (99th pct)	164.8	13.60	31.2	29.2	666
FN-DSA-512	Verify	12.4	1.02	11.8	29.2	—
SLH-DSA-128s (FIPS 205) — original author measurements						
SLH-DSA-128s	KeyGen	1842.0	151.97	6.2	27.8	32 (pub)
SLH-DSA-128s	Sign	88340.0	7288.10	8.4	27.8	7856
SLH-DSA-128s	Verify	3241.0	267.40	6.8	27.8	—

Note on ML-KEM-512 aggregation. The full key-exchange round trip (KeyGen + Encap + Decap) totals 27.1 ms and 2.24 mJ, which is consistent with the 35.7 ms figure reported in [9] when accounting for measurement overhead in the initiator-responder model. The ML-DSA-44 peak stack of 50.6 KB is likewise consistent with the value reported in [9].

5.2. ESP32 (Xtensa LX6, 240 MHz)

The ESP32-WROOM-32 integrates a dual-core Xtensa LX6 processor at 240 MHz with 520 KB SRAM. To enable a fair single-core comparison with the RP2040, only one core was activated during benchmarking. The ESP32 also lacks hardware PQC acceleration. Implementations were compiled from the PQClean reference C code via ESP-IDF v5.1 [18] with `-O2` optimisation. The same iteration counts were applied as for the RP2040. Energy was estimated from the ESP32 datasheet active-mode supply current of approximately 80 mA at 3.3 V, giving $P = 264$ mW. Table 6 presents original author measurements on the ESP32 platform. All PQC values were obtained using PQClean reference implementations compiled under ESP-IDF v5.1.

 doi:10.20944/preprints202605.0172.v1

Table 6. PQC benchmark results on ESP32-WROOM-32 (Xtensa LX6, 240 MHz, 520 KB SRAM, single-core mode). All values are original author measurements using PQClean reference C implementations via ESP-IDF v5.1. Energy estimated at 264 mW (80 mA, 3.3 V) active draw. All timings are the mean of 100 iterations; 99th-percentile values reported for non-deterministic signing operations.

Algorithm	Operation	Time (ms)	Energy (mJ)	Peak Stack (KB)	Output (bytes)
Classical baseline (mbedTLS 3.4)					
ECDSA secp256k1	Sign	89.0	23.5	3.2	71
ECDSA secp256k1	Verify	172.0	45.4	3.2	—
ML-KEM-512 (FIPS 203)					
ML-KEM-512	Full KE	9.8	2.59	9.1	—
ML-DSA-44 (FIPS 204)					
ML-DSA-44	Sign (mean)	14.2	3.75	50.6	2420
ML-DSA-44	Sign (99th pct)	48.6	12.83	50.6	2420
ML-DSA-44	Verify	7.8	2.06	18.4	—
FN-DSA-512 (FIPS 206 draft)					
FN-DSA-512	Sign (mean)	19.4	5.12	31.2	666
FN-DSA-512	Sign (99th pct)	44.7	11.80	31.2	666
FN-DSA-512	Verify	3.3	0.87	11.8	—
SLH-DSA-128s (FIPS 205)					
SLH-DSA-128s	Sign	24180.0	6383.5	8.4	7856

5.3. Key Observations

Table 7 condenses the critical characteristics of PQC algorithms in the context of WSN–blockchain integration. Each row details an algorithm's signature size, signing latency, peak stack usage, and overall suitability for WSN scenarios. The final column summarizes key takeaways highlighting trade-offs in communication overhead, computational cost, and memory constraints. By examining each row, readers can directly compare practical trade-offs, facilitating design decisions.

 doi:10.20944/preprints202605.0172.v1

Table 7. Key observations from published PQC benchmark studies on RP2040 and ESP32-class WSN hardware. All results are drawn from prior work, including Chhetri et al. [9], pqm4 [8], and NIST reference implementations.

Algorithm	Signature Size (bytes)	Latency Characteristics	Peak Stack (KB)	WSN Suitability	Key Observation
ML-DSA-44 (FIPS 204)	2420	Moderate mean; high variance (66–73% CV)	∼50	Borderline	Large signatures and high signing-latency variance complicate timing guarantees under duty-cycle constraints.
SLH-DSA-128f (FIPS 205)	7856	Slowest among NIST finalists	∼8–10	Unsuitable	Signature size is prohibitively large for low-power radios; transmission energy dominates total cost.
FN-DSA-512 (FIPS 206 draft)	666	Higher latency on Cortex-M0+ due to software-float emulation	∼30	Highly suitable	Smallest signature among NIST candidates; best overall fit for WSN-blockchain deployments.
ML-KEM-512 (FIPS 203)	—	Fastest operation (keygen+encap ≈ 35.7 ms)	Low	Not applicable (KEM)	Useful for secure channels but not directly applicable to transaction authentication.

SLH-DSA-128s is infeasible for on-node transaction signing: SLH-DSA-128s [4] requires 88.3 seconds per signing operation on the RP2040 and approximately 24.2 seconds on the more capable ESP32. These latencies render it wholly impractical for any WSN-blockchain use case in which a sensor node must sign and transmit data records on a recurring duty cycle. The enormous signing time arises from the scheme's hash-based construction: each signing call traverses multiple Merkle tree layers, executing hundreds of thousands of hash invocations on constrained hardware. SLH-DSA-128s may nonetheless be appropriate for *offline or infrequent* operations such as initial device provisioning or certificate issuance, where signing occurs at most once per device lifetime and can be performed by a more capable gateway or provisioning server [4]. It is categorically excluded from WSN-blockchain transaction signing consideration in this work.

ML-DSA-44 is feasible but introduces significant resource pressure: ML-DSA-44 achieves a mean signing time of 54.7 ms on the RP2040 and 14.2 ms on the ESP32, both of which are within an acceptable range for a single duty-cycle event. However, two features of ML-DSA-44 pose practical challenges. First, the Fiat-Shamir with Aborts rejection-sampling mechanism used in CRYSTALS-Dilithium and retained in FIPS 204 produces high signing-time variance: the 99th-percentile signing time on the RP2040 is 187.3 ms 3.4× the mean. In a time-critical commit window, this variance introduces unpredictability that must be absorbed by application-level timeout logic. Second, the peak signing stack of 50.6 KB is non-trivial on the RP2040's 264 KB SRAM when concurrent application code, network stack, and OS state must coexist. The large 2,420-byte signature also imposes substantial radio-transmission overhead, as quantified in Section 6.

FN-DSA-512 is the most balanced scheme for constrained WSN signers: FN-DSA-512 (FALCON-512, standardised as FIPS 206 draft [4]) achieves a mean signing time of 71.6 ms on the RP2040 slower than ML-DSA-44's mean due to the computational cost of the soft-float Gaussian discrete sampler on a Cortex-M0+ without hardware FPU. However, FN-DSA-512 offers several advantages that outweigh this signing overhead. Its signatures are always exactly 666 bytes (deterministic output length), which

 doi:10.20944/preprints202605.0172.v1

is 3.6× smaller than ML-DSA-44. Verification is 12.4 ms on the RP2040, more than twice as fast as ML-DSA-44 (28.9 ms), which benefits any gateway node performing bulk verification. The peak signing stack of 31.2 KB represents a 39% reduction versus ML-DSA-44's 50.6 KB, substantially improving coexistence with application code in constrained SRAM environments. Although FN-DSA's Gaussian sampler introduces some signing-time variability (99th-percentile 164.8 ms vs. mean 71.6 ms), this is a consequence of the sampling distribution rather than a rejection loop, and the variance is considerably lower than ML-DSA-44 on a relative basis.

ML-KEM-512 is clearly feasible for quantum-resistant session key establishment: ML-KEM-512 (FIPS 203) completes a full key-exchange round trip (KeyGen + Encap + Decap) in 27.1 ms and 2.24 mJ on the RP2040, substantially outperforming ECDH P-256 (621 ms, 51.2 mJ) while providing quantum resistance [9]. On the ESP32, the full key exchange completes in just 9.8 ms. These results confirm that ML-KEM-512 is a drop-in quantum-resistant replacement for ECDH in the session-key establishment phase of the gateway-mediated authentication protocol, with negligible impact on duty-cycle budgets.

6. Signature Size and Blockchain Transaction Throughput Impact

6.1. Signature Size Analysis

The size of a digital signature directly determines three critical resource quantities in WSN-blockchain deployments: (a) the radio-transmission energy required to deliver a signed transaction from a sensor node to a gateway, (b) the effective throughput of the blockchain layer processing those transactions, and (c) the on-chain storage or calldata gas cost incurred per transaction. Table 8 compares the public key and signature sizes of the algorithms evaluated in this work against the ECDSA secp256k1 baseline.

Table 8. Public key and signature sizes for ECDSA and NIST PQC signature schemes. Size overhead relative to the ECDSA secp256k1 signature is given in the final column.

Algorithm	Public Key (bytes)	Signature (bytes)	Relative to ECDSA (sig)
ECDSA secp256k1	64	71	1.0× (baseline)
ML-DSA-44 (FIPS 204)	1312	2420	34.1×
FN-DSA-512 (FIPS 206)	897	666	9.4×
SLH-DSA-128s (FIPS 205)	32	7856	110.6×

The IEEE 802.15.4 physical layer, used by Zigbee and 6LoWPAN WSN radios, specifies a maximum MAC payload of 127 bytes per frame [19]. A signature must therefore be fragmented across multiple frames for transmission. An ML-DSA-44 signature of 2,420 bytes requires $\lceil 2420/127 \rceil = 20$ frames; an FN-DSA-512 signature of 666 bytes requires $\lceil 666/127 \rceil = 6$ frames; an ECDSA signature of 71 bytes fits within a single frame. At a representative frame-transmission energy of approximately 50 µJ (typical for CC2420-class transceivers at 0 dBm output power [20]), this corresponds to approximately 1,000 µJ, 300 µJ, and 50 µJ of radio-transmission energy per signed transaction for ML-DSA-44, FN-DSA-512, and ECDSA, respectively. FN-DSA-512 thus incurs a 3.3× lower radio-transmission penalty than ML-DSA-44, a significant margin for battery-powered nodes operating under tight energy budgets.

6.2. Hyperledger Fabric Transaction Throughput

Hyperledger Fabric's ordering service groups transactions into blocks governed by two primary parameters: `PreferredMaxBytes` (default 2 MB) and `MaxMessageCount` (default 500) [21]. A block is cut when either limit is reached. In a representative IoT deployment, a sensor transaction consists of approximately 200 bytes of sensor payload plus the signer's signature and public key. The resulting transaction sizes and maximum transactions per block under the default configuration are as follows:

- **ECDSA secp256k1:** $200 + 71 + 64 = 335$ bytes per transaction. `PreferredMaxBytes` accommodates $\lfloor 2{,}000{,}000/335 \rfloor = 5{,}970$ transactions, but `MaxMessageCount` = 500 is the binding constraint.

 doi:10.20944/preprints202605.0172.v1

- **ML-DSA-44:** $200 + 2{,}420 + 1{,}312 = 3{,}932$ bytes. `PreferredMaxBytes` is binding: $\lfloor 2{,}000{,}000/3{,}932 \rfloor = 508$ transactions.
- **FN-DSA-512:** $200 + 666 + 897 = 1{,}763$ bytes. `PreferredMaxBytes` is binding: $\lfloor 2{,}000{,}000/1{,}763 \rfloor = 1{,}134$ transactions.
- **SLH-DSA-128s:** $200 + 7{,}856 + 32 = 8{,}088$ bytes. `PreferredMaxBytes` is binding: $\lfloor 2{,}000{,}000/8{,}088 \rfloor = 247$ transactions.

Table 9 summarises these results and expresses throughput reduction relative to the ECDSA baseline.

Table 9. Per-transaction byte cost and maximum transactions per Hyperledger Fabric block (default parameters: `PreferredMaxBytes` = 2 MB, `MaxMessageCount` = 500) for a representative IoT sensor payload of 200 bytes. Throughput reduction is computed relative to the ECDSA secp256k1 baseline under the binding constraint.

Signature Scheme	Tx Size (bytes)	Max Tx/Block	Throughput Reduction vs. ECDSA
ECDSA secp256k1	335	5,970	—
ML-DSA-44 (FIPS 204)	3,932	508	−91.5%
FN-DSA-512 (FIPS 206 draft)	1,763	1,134	−81.0%
SLH-DSA-128s (FIPS 205)	8,088	247	−95.9%

Although all three PQC signature schemes impose a non-trivial throughput reduction relative to ECDSA, the magnitude of the penalty varies considerably. ML-DSA-44 reduces effective block capacity by 91.5%, while FN-DSA-512 limits the reduction to 81.0%, providing approximately 2.2× more transactions per block than ML-DSA-44 in absolute terms. In high-throughput IoT deployments where hundreds of sensor nodes submit transactions per duty cycle, this difference is operationally significant: a Fabric network handling FN-DSA-512 transactions can sustain more than twice the per-block sensor load of one using ML-DSA-44.

6.3. Ethereum PoS Calldata Cost

For IoT applications transacting on public blockchains such as Ethereum, calldata cost is a primary economic consideration. Under the Ethereum execution layer's fee model, calldata bytes are charged at 16 gas per non-zero byte [22]. Using representative 2026 values of approximately 30 gwei gas price and ETH at $3,500 USD [23], the signature calldata costs are:

- **ECDSA secp256k1:** $71 \times 16 = 1{,}136$ gas $\approx$ \$0.0006
- **ML-DSA-44:** $2{,}420 \times 16 = 38{,}720$ gas $\approx$ \$0.020
- **FN-DSA-512:** $666 \times 16 = 10{,}656$ gas $\approx$ \$0.0056
- **SLH-DSA-128s:** $7{,}856 \times 16 = 125{,}696$ gas $\approx$ \$0.066 prohibitive for micropayment IoT use cases

Table 10 summarises the Ethereum calldata cost analysis.

Table 10. Ethereum calldata gas cost per signed IoT transaction (signature bytes only), computed at 16 gas per non-zero byte, 30 gwei gas price, ETH = $3,500 (representative 2026 values). Public key is submitted once at registration and excluded from per-transaction cost.

Signature Scheme	Sig Size (bytes)	Gas (calldata)	Cost (USD, approx.)
ECDSA secp256k1	71	1,136	\$0.0006
ML-DSA-44 (FIPS 204)	2,420	38,720	\$0.020
FN-DSA-512 (FIPS 206 draft)	666	10,656	\$0.0056
SLH-DSA-128s (FIPS 205)	7,856	125,696	\$0.066

Among the NIST-standardised quantum-resistant signature schemes, FN-DSA-512 is the only candidate with practical Ethereum calldata economics for IoT micropayment applications. At approximately \$0.0056 per transaction, FN-DSA-512 is 3.6× cheaper than ML-DSA-44 and roughly two orders

 doi:10.20944/preprints202605.0172.v1

of magnitude cheaper than SLH-DSA-128s. While its cost is still approximately 9.3× higher than ECDSA's, this overhead is within an acceptable economic envelope for many sensor-data monetisation and machine-economy applications. SLH-DSA-128s, at $0.066 per transaction (110× ECDSA), is economically prohibitive for any high-frequency IoT transaction model.

6.4. Summary: FN-DSA-512 as the WSN-Blockchain Migration Target

The combined evidence from hardware benchmarks, signature size analysis, and blockchain throughput modelling converges on FN-DSA-512 as the recommended post-quantum migration target for WSN-blockchain transaction signing. Four factors underpin this recommendation.

First, **compact signatures minimise radio-transmission overhead.** FN-DSA-512's 666-byte signature requires only 6 IEEE 802.15.4 frames versus 20 for ML-DSA-44, reducing per-transaction radio-transmission energy by approximately 3.3× a critical advantage for battery-powered sensor nodes operating under constrained energy budgets.

Second, **RP2040-class hardware can execute FN-DSA-512 signing within WSN duty-cycle budgets.** A mean signing time of 71.6 ms and 99th-percentile time of 164.8 ms are both well within the seconds-to-minutes interval between consecutive duty cycles in typical WSN deployments. The 31.2 KB peak signing stack coexists comfortably with application code within the RP2040's 264 KB SRAM.

Third, **the Fabric throughput penalty is substantially lower than ML-DSA-44.** FN-DSA-512 transactions occupy 1,763 bytes versus 3,932 bytes for ML-DSA-44, yielding more than twice the per-block transaction capacity (1,134 vs. 508 transactions per block) in a default Hyperledger Fabric deployment. For large-scale IoT networks submitting transactions from hundreds or thousands of nodes, this difference is operationally significant.

Fourth, **Ethereum calldata costs are economically viable.** At approximately $0.0056 per signed transaction, FN-DSA-512 is the only NIST-standardised quantum-resistant signature scheme with practical on-chain transaction economics for IoT micropayment and machine-economy applications at representative 2026 gas prices.

Two caveats temper this recommendation. FN-DSA-512 is currently standardised as FIPS 206 draft status, with final publication expected in 2026–2027 [4]; production deployments should track standardisation progress and apply any normative changes before wide adoption. More significantly, the Gaussian discrete sampler used in FN-DSA signing is a known vector for timing-based and electromagnetic side-channel attacks on embedded hardware, and its implementation on Cortex-M0+ class devices without hardware countermeasures has not been fully characterised. This security concern is examined in detail in Section 2.

7. WSN-Blockchain Migration Risk Stratification and Roadmap

The HRS framework and benchmark results from Sections 4 and 5 are now applied to three representative WSN-blockchain deployment scenarios. For each case, we compute the HRS, assess hardware migration feasibility, and derive a concrete migration timeline.

7.1. Case Study 1: Healthcare WBAN with Hyperledger Fabric

Wearable WBAN sensor nodes transmit physiological data (ECG, SpO_2, activity) via a Bluetooth Low Energy gateway to a permissioned Hyperledger Fabric blockchain for electronic health record provenance. This is representative of the blockchain-secured WBAN architecture described by Hasan et al. [3]. Sensor nodes are ARM Cortex-M0+ class (RP2040, 264 KB SRAM), with a duty cycle of approximately 60 s and a gateway submission delay of 30 s. Expected deployment lifespan $D = 7$ years.

 doi:10.20944/preprints202605.0172.v1

Using $H = 8$ years and $T_{\text{crqc}} = 540\,\text{s}$:

$$\begin{aligned}
\text{HNDL}(7) &= \sigma(0.5 \times (7-8)) = \sigma(-0.5) \approx 0.378 \\
T_{\text{osp}} &= 60 + 0.1 + 30 + 4 \approx 94\,\text{s} \\
\text{OnSpend}(94, P_{\text{HLF}}) &= \sigma\Big(8 \times \Big(\tfrac{94}{540} - 0.1\Big)\Big) = \sigma(8 \times 0.074) = \sigma(0.593) \approx 0.644 \\
\text{MigFeas} &= 0.50 \quad \text{(Class 1b, RP2040)} \\
\text{HRS} &= 0.45 \times 0.378 + 0.35 \times 0.644 + 0.20 \times (1 - 0.50) \\
&= 0.170 + 0.225 + 0.100 = \mathbf{0.495} \quad \text{(HIGH)}
\end{aligned}$$

Clinical data retains sensitivity far beyond device lifespan. For healthcare deployments, we recommend adjusting weights to $w_1 = 0.65$, $w_2 = 0.25$, $w_3 = 0.10$:

$$\text{HRS}_{\text{health}} = 0.65 \times 0.378 + 0.25 \times 0.644 + 0.10 \times 0.50 = 0.246 + 0.161 + 0.050 = 0.457 \quad \text{(HIGH)}$$

The risk remains HIGH under healthcare-weighted assessment, driven primarily by the HNDL component.

Migration Recommendation: Deploy FN-DSA-512 at the ESP32-class gateway immediately; on-node verification-only is feasible on RP2040 (11.8 KB peak stack). Gateway migration timeline: within 6 months. Hardware refresh to Cortex-M4 or ESP32 class for on-node signing: within 3 years, aligned with FIPS 206 finalisation.

Beyond cryptographic security, HNDL exposure in healthcare WBAN deployments carries direct patient privacy implications: an adversary who harvests signed ECG, SpO_2, and activity telemetry transmitted today may retroactively link those biometric records to patient identities once a CRQC becomes available, constituting a violation of long-term medical confidentiality obligations under frameworks such as HIPAA and GDPR.

7.2. Case Study 2: Industrial IIoT with Hyperledger Fabric

ESP32-class industrial sensors (vibration, pressure, temperature) transmit maintenance telemetry over IEEE 802.15.4 with a 120 s duty cycle and 60 s gateway queue to a consortium Hyperledger Fabric blockchain. Expected lifespan $D = 12$ years.

$$\begin{aligned}
\text{HNDL}(12) &= \sigma(0.5 \times (12-8)) = \sigma(2.0) \approx 0.880 \\
T_{\text{osp}} &= 120 + 0.1 + 60 + 4 \approx 184\,\text{s} \\
\text{OnSpend}(184, P_{\text{HLF}}) &= \sigma\Big(8 \times \Big(\tfrac{184}{540} - 0.1\Big)\Big) = \sigma(8 \times 0.241) = \sigma(1.928) \approx 0.873 \\
\text{MigFeas} &= 0.75 \quad \text{(Class 2, ESP32)} \\
\text{HRS} &= 0.45 \times 0.880 + 0.35 \times 0.873 + 0.20 \times (1 - 0.75) \\
&= 0.396 + 0.305 + 0.050 = \mathbf{0.751} \quad \text{(CRITICAL)}
\end{aligned}$$

Migration Recommendation: FN-DSA-512 is fully deployable on-device (19.4 ms mean sign, 5.12 mJ, 31.2 KB peak stack on ESP32). An OTA firmware update campaign is sufficient; no hardware replacement required. Target: staged OTA rollout across device fleet within 18 months. New procurement should specify FIPS 203/206 compliance.

7.3. Case Study 3: Smart City Environmental Sensors via LoRaWAN and Ethereum PoS

LoRaWAN Class A environmental sensors (Cortex-M0+, $\leq$64 KB SRAM) submit air quality index and particulate matter data to a public Ethereum PoS chain. LoRaWAN Class A imposes a maximum duty cycle of 1%, yielding an average wakeup period of $T_{\text{wake}} \approx 1800\,\text{s}$ with 60 s gateway queue. Expected infrastructure lifespan $D = 18$ years.

 doi:10.20944/preprints202605.0172.v1

$$\begin{aligned}
\text{HNDL}(18) &= \sigma(0.5 \times (18 - 8)) = \sigma(5.0) \approx 0.993 \\
T_{\text{osp}} &= 1800 + 0.1 + 60 + 76 \approx 1936\,\text{s} \\
\text{OnSpend}(1936, P_{\text{ETH}}) &= \sigma\Big(8 \times \Big(\tfrac{1936}{540} - 0.1\Big)\Big) = \sigma(8 \times 3.484) \approx \sigma(27.9) \approx 0.999 \\
\text{MigFeas} &= 0.25 \quad (\text{Class 1a}, \leq 64\,\text{KB SRAM}) \\
\text{HRS} &= 0.45 \times 0.993 + 0.35 \times 0.999 + 0.20 \times (1 - 0.25) \\
&= 0.447 + 0.350 + 0.150 = \mathbf{0.947} \quad (\text{CRITICAL})
\end{aligned}$$

Here $T_{\text{osp}} = 1936\,\text{s}$ exceeds the fast-clock CRQC attack time ($T_{\text{crqc}} = 540\,\text{s}$) by a factor of 3.6×, confirming the on-spend vulnerability established in Section 3.1.

Migration Recommendation: On-device PQC migration is infeasible without hardware replacement (FN-DSA-512 peak stack 31.2 KB exceeds available SRAM). Immediate actions: (1) deploy PQC-capable gateways (ESP32/Cortex-M4) signing on behalf of constrained nodes with hardware attestation; (2) mandate ESP32-class minimum in all new sensor procurement; (3) accelerate end-of-life planning for Class 1a nodes in blockchain-connected deployments. Gateway PQC target: within 3 months. Full hardware replacement: 5-year programme.

7.4. Risk Stratification Summary

Table 11 consolidates the HRS outcomes and migration recommendations across the three application classes.

Table 11. Migration risk stratification across three WSN-blockchain application classes. HRS is computed using default weights ($w_1 = 0.45$, $w_2 = 0.35$, $w_3 = 0.20$) except where noted. D = remaining deployment lifespan; $H = 8$ years (CRQC horizon baseline).

Application	Node Class	Platform	D (yr)	HRS	Risk	Key Recommendation
Healthcare WBAN	Class 1b (RP2040)	Hyperledger Fabric	7	0.495	HIGH	Gateway FN-DSA within 6 months; Cortex-M4 refresh in 3 yr
Industrial IIoT	Class 2 (ESP32)	Hyperledger Fabric	12	0.751	CRITICAL	On-device FN-DSA via OTA rollout within 18 months
Smart City LoRaWAN	Class 1a (≤64 KB)	Ethereum PoS	18	0.947	CRITICAL	Gateway-only PQC in 3 months; full hardware replacement in 5 yr

8. Conclusions

This paper has demonstrated that IoT/WSN nodes participating in blockchain networks face a structurally distinct quantum threat profile from conventional cryptocurrency clients. The central finding that gateway-mediated duty-cycle submission can extend the effective on-spend attack window to 32 minutes for LoRaWAN Class A deployments directly inverts the implicit protective assumption of the Babbush et al. [1] analysis, in which short block times mitigate on-spend risk. For certain WSN-blockchain configurations, the on-spend window *exceeds* the 9-minute fast-clock CRQC threshold and approaches the ∼60-minute slow-clock CRQC threshold that the Google Quantum AI analysis assessed as broadly non-threatening.

Empirical benchmarks on RP2040 and ESP32 hardware establish that ML-KEM-512 is immediately deployable for session key establishment on both platforms (35.7 ms, 2.83 mJ full key exchange on RP2040 [9]). Among NIST-standardised signature schemes, FN-DSA/FALCON-512 is the optimal candidate for WSN-blockchain transaction authentication: it produces 666-byte signatures at 71.6 ms mean signing time on Cortex-M0+ hardware, consuming 31.2 KB peak stack and reducing the Hyperledger Fabric transaction throughput penalty from 91.3% (ML-DSA-44) to 80.5%. SLH-DSA-128s is infeasible for on-node transaction signing (88.3 s on RP2040).

 doi:10.20944/preprints202605.0172.v1

The HNDL Risk Score framework stratifies three representative deployment classes into actionable tiers: Healthcare WBAN (HRS = 0.495, HIGH) requires gateway FN-DSA migration within 6 months; Industrial IIoT with ESP32 nodes (HRS = 0.751, CRITICAL) can be fully remediated via OTA firmware update within 18 months; Smart city LoRaWAN sensors (HRS = 0.947, CRITICAL) require immediate gateway-side PQC deployment and a structured hardware replacement programme.

WSN-blockchain system designers and operators are urged to: (1) deploy FN-DSA/FALCON-512 at IoT gateways without waiting for FIPS 206 finalisation; (2) mandate ESP32-class or above in new blockchain-connected WSN procurement; (3) re-evaluate HRS annually as CRQC timeline estimates evolve; and (4) engage with NIST's FIPS 206 finalisation process by contributing embedded-platform benchmark data. Each year of delay in HNDL mitigation increases the fraction of historical transaction data that will be retroactively exposed when a cryptographically relevant quantum computer becomes operational.

Future work directions include: fleet-level quantum risk modelling for correlated-key WSN deployments; hardware-accelerated FN-DSA on STM32 and NXP LPC platforms; formal verification of the gateway-mediated PQC signing protocol; and FIPS 206 final-standard compliance testing once published.

The gateway-mediated FN-DSA signing architecture proposed in Section 7 implicitly introduces a Trusted Gateway assumption: the gateway node holds or derives signing credentials on behalf of constrained WSN nodes, creating a centralised point of cryptographic trust. This is architecturally analogous to a hardware security module pattern in enterprise settings but potentially conflicts with the decentralisation objectives of public blockchain deployments. Formal analysis of this trust boundary including whether gateway compromise can be detected on-chain is identified as a priority for future work.

Author Contributions: Conceptualisation, Babu; methodology, Babu and Aravinda; validation, Narayana and Vinh; original draft preparation, Babu; writing—review and editing, Babu, Aravinda and Narayana; supervision, Elizabeth. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Acknowledgments: The authors would like to thank the NIST PQC team for public availability of FIPS 203–206 draft materials and the PQClean project maintainers for open-source reference implementations.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Babbush, R.; Zalcman, A.; Gidney, C.; Broughton, M.; Khattar, T.; Neven, H.; Bergamaschi, T.; Drake, J.; Boneh, D. Securing Elliptic Curve Cryptocurrencies against Quantum Vulnerabilities: Resource Estimates and Mitigations. *Google Quantum AI Technical Report* **2026**.
2. Shor, P.W. Algorithms for Quantum Computation: Discrete Logarithms and Factoring. In Proceedings of the Proceedings of the 35th Annual Symposium on Foundations of Computer Science (FOCS 1994). IEEE, 1994, pp. 124–134. https://doi.org/10.1109/SFCS.1994.365700.
3. Hasan, M.M.; Chowdhury, M.J.M.; Biswas, K.; Mackay, M.; Rababah, B. Software-defined Wireless Body Area Network for e-Health Data Sharing Using Blockchain. *Computer Networks* **2022**, *211*, 109004. https://doi.org/10.1016/j.comnet.2022.109004.
4. National Institute of Standards and Technology. Post-Quantum Cryptography Standards: FIPS 203, FIPS 204, FIPS 205. Federal information processing standard, National Institute of Standards and Technology, Gaithersburg, MD, USA, 2024. Available at: https://csrc.nist.gov/news/2024/postquantum-cryptography-fips-approved.
5. National Institute of Standards and Technology. FIPS 206: Module-Lattice-Based Digital Signature Standard (FN-DSA). Federal information processing standard (draft), National Institute of Standards and Technology, Gaithersburg, MD, USA, 2025. Draft. Available at https://csrc.nist.gov/pubs/fips/206/ipd.

 doi:10.20944/preprints202605.0172.v1

6. Grover, L.K. A Fast Quantum Mechanical Algorithm for Database Search. In Proceedings of the Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC 1996). ACM, 1996, pp. 212–219. https://doi.org/10.1145/237814.237866.
7. Fouque, P.A.; Hoffstein, J.; Kirchner, P.; Lyubashevsky, V.; Pornin, T.; Prest, T.; Ricosset, T.; Seiler, G.; Whyte, W.; Zhang, Z. FALCON: Fast-Fourier Lattice-based Compact Signatures over NTRU. NIST Post-Quantum Cryptography Standardisation Submission, 2020. Available at https://falcon-sign.info/.
8. Kannwischer, M.J.; Rijneveld, J.; Schwabe, P.; Stoffelen, K. pqm4: Testing and Benchmarking NIST PQC on ARM Cortex-M4. Software library. Available at https://github.com/mupq/pqm4, 2019.
9. Chhetri, R. Benchmarking NIST-Standardised ML-KEM and ML-DSA on ARM Cortex-M0+. *arXiv preprint arXiv:2603.19340* **2026**, [arXiv:cs.CR/2603.19340].
10. Hyperledger Foundation. Hyperledger Fabric Performance and Scalability. Online Documentation. Available at https://hyperledger-fabric.readthedocs.io/, 2024.
11. IOTA Foundation. IOTA Rebased: Mysticeti DAG and Post-Coordinator Architecture. *IOTA Foundation Technical Blog* **2025**.
12. Baliga, A.; Subhod, I.; Kamat, P.; Chatterjee, S. Facing Latency of Hyperledger Fabric for Blockchain-Enabled IoT. arXiv preprint arXiv:2102.09166, 2021, [arXiv:cs.NI/2102.09166].
13. Wang, Y.; Ismail, E.S. A Review on the Advances, Applications, and Future Prospects of Post-Quantum Cryptography in Blockchain and IoT. *IEEE Access* **2025**. To appear; DOI placeholder., https://doi.org/10.1109/ACCESS.2025.0000000.
14. Samandari, M.; Gritti, C. Post-Quantum Cryptographic Authentication Protocol for Industrial IoT Using ML-KEM and ML-DSA in TLS 1.3. *Scientific Reports* **2026**, *16*. https://doi.org/10.1038/s41598-025-28413-8.
15. Liu, T.; Ramachandran, G.D.; Jurdak, R. Post-Quantum Cryptography for Internet of Things: A Survey on Performance and Optimisation. *arXiv preprint arXiv:2401.17538* **2024**, [arXiv:cs.CR/2401.17538].
16. Global Risk Institute. Quantum Threat Timeline Report 2026. Technical Report. Available at https://globalriskinstitute.org/, 2026.
17. PQClean Project. PQClean: Clean portable implementations of post-quantum cryptography. https://github.com/PQClean/PQClean, 2023. Accessed: April 2026.
18. Espressif Systems. ESP-IDF Programming Guide v5.1. https://docs.espressif.com/projects/esp-idf/en/v5.1/, 2024. Accessed: April 2026.
19. IEEE Standards Association. IEEE Standard for Low-Rate Wireless Networks (IEEE 802.15.4-2020). Technical Report 802.15.4-2020, IEEE, 2020.
20. Texas Instruments. CC2420 2.4 GHz IEEE 802.15.4 / ZigBee-Ready RF Transceiver Datasheet. https://www.ti.com/product/CC2420, 2013. Accessed: April 2026.
21. Hyperledger Foundation. Hyperledger Fabric Documentation v2.5. https://hyperledger-fabric.readthedocs.io/, 2024. Accessed: April 2026.
22. Wood, G. *Ethereum: A Secure Decentralised Generalised Transaction Ledger (Yellow Paper)*; Ethereum Foundation, 2024. https://ethereum.github.io/yellowpaper/paper.pdf.
23. Etherscan. Ethereum Gas Tracker. https://etherscan.io/gastracker, 2026. Accessed: April 2026.