P3Q Pulsar Kind: Module-LWE Threshold Verify on the Unified Precompile Slot

Zach Kelling

Lux Industries

2026-06-03

Abstract

This paper specifies *kind* $0x01$ *Pulsar* of the P3Q framework (LP-219): a rank-k Module-LWE threshold verifier over FIPS-204 ML-DSA-65. Pulsar is the lowest-gas P3Q kind (50 000 gas) and the first wired on-chain; it is the PQ-fast leg of the LP-217 cert profile ladder. The slot, the kind-byte dispatch ABI, and the `P3QVerifier` registration interface are specified in LP-219; this paper gives Pulsar's wire layout, its parameter set, and its gas calibration. Pulsar reduces to the hardness of rank-k Module-LWE, the assumption underlying the FIPS-204 standard, and is independent of the Module-LWE (Corona, LP-220) and hash-collision (Magnetar, LP-222) assumptions its sibling kinds rest on. The CPU verifier is byte-equal to the FIPS-204 standard and is wired today; GPU dispatch is projected, not measured.

Contents

1 Introduction

P3Q (LP-219) is the unified-slot post-quantum verifier framework at EVM precompile slot `0x012205`: one slot, one ABI, and a leading `kind` byte that dispatches into a per-kind threshold-signature verifier core. This paper specifies *kind* $0x01$ *Pulsar*, the Module-LWE member of the P3Q family.

For the slot allocation, the kind-byte dispatch ABI, the `P3QVerifier` Go interface that registers each kind, and the cross-family composition into LP-217 cert profile modes, see LP-219. The sibling kinds are specified in LP-220 (Corona / Module-LWE) and LP-222 (Magnetar / hash-based FIPS-205 SLH-DSA).

Pulsar reduces to the hardness of Module-LWE — the assumption underlying the FIPS-204 ML-DSA standard — which it shares with its sibling Corona (Module-LWE at the Ringtail parameter set); the two differ in parameter regime and codebase, not hardness family. Pulsar's family-independent sibling is the hash-based Magnetar kind, whose SHA-3 collision assumption is disjoint from the lattice family. It is the lowest-gas P3Q kind (50 000 gas), the first wired on-chain, and the PQ-fast leg of the LP-217 cert profile ladder. Its CPU verifier is byte-equal to the FIPS-204 standard and is wired today; throughout this paper performance numbers are tagged **measured** or **projected**, and GPU dispatch is projected, not measured.

2 Pulsar Kind (kind = $0x01$)

2.1 Slot and ABI

`0x012205` with leading `kind=0x01`. The Solidity shim is:

```
// EVM precompile at slot 0x012205, kind = 0x01
bool ok = P3Q.verify(
    0x01,                       // Pulsar kind
    pulsarSig,                  // FIPS-204 ML-DSA-{44|65|87} sig
    messageHash,                // H(prev_root || new_root || batch_hash)
    groupPubKey                 // Pulsar group public key
);
```

The kind byte at offset 0 selects dispatch into the Pulsar verifier; the mode byte at offset 1 selects which FIPS-204 parameter set is in play.

2.2 Wire format

The full 5303 B canonical input for ML-DSA-65 (the production target):

Offset	Size	Field
0	1	`kind = 0x01`
1	1	`mode` (`0x44`/`0x65`/`0x87` = ML-DSA-44/65/87)
2	4	`sigLen` (BE32)
6	`sigLen`	`pulsarSig` (3309 B for ML-DSA-65)
6+*sigLen*	4	`pkLen` (BE32)
10+*sigLen*	`pkLen`	`groupPubKey` (1952 B for ML-DSA-65)
10+*sigLen*+*pkLen*	32	`messageHash`

Table 1: Pulsar-kind wire format. Total at ML-DSA-65: 1+1+4+3309+4+1952+32 = 5303 B.

Trailing bytes after the messageHash are accepted as padding to match the convention used by the sibling STARK-FRI precompile; padding does not affect verification.

2.3 Verifier core

The Pulsar verifier core IS the FIPS-204 ML-DSA verifier. Pulsar signatures serialize byte-identically to FIPS-204 ML-DSA signatures per LP-073 [1] – the Pulsar threshold ceremony's Class-N1 claim is that the aggregated signature is byte-equal to a single-party FIPS-204 ML-DSA signature on the same message and group public key. Any conformant FIPS-204 verifier suffices; the reference implementation at `lux/precompile/p3q/contract.go::Run` dispatches via `mldsa.PublicKeyFromBytes(...) + pk.VerifySignatureCtx(...)`, which is a thin alias over the cloudflare/circl FIPS-204 implementation.

The Pulsar kind is the only kind with a fully-wired verifier today. Corona ($0x02$) and Magnetar ($0x03$) reserve their dispatch entries and return `abiFalse` via the `ErrKindNotImplemented` path until their threshold libraries land. See LP-220 and LP-222.

2.4 Gas cost

50 000 gas base + 10 gas/byte per-byte adder. At the canonical 5303 B input, total gas is $50000 + 5303 \times 10 = 103\,030$ gas. This is well under the 12 M C-Chain block-gas-limit (~116 verifies per block; per the audit at `contract.FeeConfigReporter`). The gas calibration is 10× the BLS pairing verify charge of 5000 gas, reflecting lattice verification's inherent cost (NTT inverse + matrix-vector multiplication mod q).

2.5 Domain separation

Every Pulsar-kind verify binds the FIPS-204 context string `lux-evm-precompile-p3q-v1`. Mismatched context between signer and verifier causes the FIPS-204 verifier to reject. This prevents cross-precompile replay: a Pulsar signature minted for the LP-0120 raw-Pulsar precompile at `0x012204` (which uses context `lux-evm-precompile-pulsar-v1`) cannot be re-submitted to the rollup-batch P3Q path at `0x012205`.

2.6 Performance

Metric	M1 Max	Spark	Source
Per-validator Pulsar sign at N=64	391 ms	585 ms	matrix bench (measured) [2]
CPU verify (single sig)	~80 µs	~80 µs	circl benchmark (measured)
Blackwell sm_120 verify	—	~50 µs	projected (LP-203 kernel)
Blackwell sm_120 batched (N=64)	—	~5 ms aggregate	projected

Table 2: Pulsar kind performance. Sign times are measured per the Quasar 4-scheme matrix bench (mean of 50 rounds at N=64); verify times tagged "projected" have not been exercised in the matrix bench (GPU dispatch boundary gap, see section 2.6.1).

2.6.1 Note on GPU verify projection

The Blackwell verify numbers cited in earlier LP-203 drafts (~50 µs single, ~5 ms batched) are kernel-level micro-benchmarks, not end-to-end measurements through the precompile entry-point. The matrix bench captured 4595 `nvidia-smi dmon` samples during the 2-hour Spark run: mean `SM%` = 11.9 % (idle/background only). Until `lux-lattice.pc` is installed on Spark and `lux_pulsar_verify_batch` is exposed in `lux/accel/c_api.h`, GPU verify timings remain projections.

2.7 Use case

Pulsar-kind is the default PQ-fast verifier for any rollup that wants Module-LWE-grade PQ assurance on its batch envelope without paying the cross-family premium of Corona or Magnetar. Per LP-217, PQ-fast (BLS + Pulsar) is the recommended mode for payment rails, gaming, and DEX matching. The gas cost is ∼55 000 per batch (BLS + Pulsar), well below the ∼5 M cost of adding Corona.

LP-218 specifies the Pulsar kind in detail as part of the rollup-batch envelope pattern; this section restates the Pulsar-kind facts in the LP-220 per-kind structure so the three kinds are commensurable in one document.

References

[1] Lux Industries. Lp-073: Pulsar two-round lattice threshold signing, 2026.

[2] Lux Industries. Quasar 4-Scheme Matrix Bench – m1 max + dgx spark gb10, 2026. `lux/threshold/protocols/parity/parity_test.go`, run 2026-06-04. Mean of N=$\{16, 32, 64\}$ across 50 rounds.