PORTALZ

MASTER CRYPTANALYTIC PAPER

AES Key Extraction
AN NIST Blind Spot

Two Tautological Threat Classes
both outside the Global Cryptographic Standards Framework

Wolfgang Flatow | PORTALZ PTY LTD, Australia | 2026

Master paper integrating: NNKE — Neural Network Key Extraction | QKE — Quantum Key Extraction
Consequences are detailed in companion paper: Cryptographic Armageddon

The NIST Post-Quantum Encryption (PQE) programme is the global framework for cryptographic security in the post-classical computing era. It addresses two quantum threats: Shor's algorithm against asymmetric cryptography, and Grover's algorithm against symmetric encryption. For AES specifically, NIST's position is that key doubling to AES-256 provides sufficient quantum resistance. This paper demonstrates that this position contains two critical blind spots. First, Quantum Key Extraction (QKE) is a quantum attack class categorically distinct from Grover that requires no iteration, no amplitude amplification, and no quantum error correction — and is therefore not addressed by PQE. Second, Neural Network Key Extraction (NNKE) exploits geometric key signatures in AES ciphertext using commodity AI hardware, requiring no quantum computing at all — and no NIST category exists for this threat. Both attack classes rest on logical tautologies. Both are accessible with hardware that exists today. NIST's quantum threat timeline of 10-20 years is shown to be based on the wrong threat model. We propose a new mandatory cryptographic standards category — Post Neural Net Encryption (PNNE) — and demonstrate that the Fractal Encryption Standard (FES) is the only current architecture that addresses all identified threat classes.

Contents

1. The Framework Gap

The NIST Post-Quantum Encryption programme has produced four standardised algorithms addressing the quantum threat to asymmetric cryptography: CRYSTALS-Kyber, CRYSTALS-Dilithium, FALCON, and SPHINCS+. For symmetric encryption, NIST's guidance is simpler: double the key length. AES-256 is considered quantum-resistant because Grover's algorithm reduces its effective security to 128 bits — still considered computationally infeasible.

This framework is correct and complete for the threats it considers. It is incomplete as a general security standard because it considers only two quantum attack classes and zero neural network attack classes.

> ***QKE is not in PQE. PNNE does not exist. These are not oversights of detail — they are categorical absences in the global cryptographic standards framework.***

Threat Class	NIST Framework	Coverage	Status
Shor — asymmetric crypto	PQE Programme	Full — CRYSTALS, FALCON, SPHINCS+	Addressed
Grover — AES key search	PQE / key doubling	Partial — AES-256 recommended	Partially addressed
QKE — quantum superposition + SOD	None	None — not recognized as distinct from Grover	BLIND SPOT
NNKE — neural network extraction	None	None — no category exists	BLIND SPOT

This paper addresses both absences. Section 2 introduces the Flatow Cryptanalytic Framework — the theoretical foundation common to both threat classes. Sections 3 and 4 summarise NNKE and QKE respectively. Section 5 analyses NIST's position against both. Section 6 proposes the missing framework category PNNE. Section 7 demonstrates FES as the architectural resolution.

2. The Flatow Cryptanalytic Framework

Both NNKE and QKE derive from a single foundational tautology:

> ***"A key applied directly to a payload with a fixed algorithm must produce a fixed signature in the ciphertext."***

This is not a hypothesis. It is a logical necessity. AES is a deterministic bijection. Every key K defines a unique, complete permutation of the 2^{128} plaintext space. That permutation is the signature. It exists by definition. Every ciphertext block is a sample from it.

The Flatow Cryptanalytic Framework observes that this tautology opens two distinct attack surfaces — one exploitable by neural networks, one by quantum computers — neither of which was contemplated in AES's original security proofs or in NIST's current framework.

2.1 The Burden of Proof Inverts

In conventional cryptanalysis the attacker must demonstrate weakness. The tautological foundation inverts this:

```
Classical position:   Signature exists but extraction is computationally
                      infeasible by algebraic and statistical tools

Flatow position:      The extraction tool class (NN, QC) was never
                      included in the original hardness assumption.
                      The burden of proof lies with those who claim
                      AES is safe against these classes.
                      That proof does not exist.
```

2.2 Two Attack Surfaces, One Root Cause

The root cause of both vulnerabilities is identical: AES's fixed transformation per key. The attack surfaces differ only in the tool used to exploit it:

```
NNKE:   Fixed permutation → geometric invariant in 128-dimensional space
         → detectable by neural network as universal function approximator
         → no quantum hardware required
         → accessible TODAY with commodity GPUs

QKE:    Fixed permutation → unique non-random plaintext per key
         → deterministic SOD gate in quantum superposition
         → single-cycle key recovery without iteration
         → accessible NOW with Osprey-class hardware
```

3. Neural Network Key Extraction (NNKE)

Full treatment: "Neural Net Key Extraction (NNKE): A Tautological Argument for AES Neural Net Vulnerability" (Flatow, 2026).

3.1 The Core Argument

AES with a fixed key is a deterministic bijection. Its unique permutation of the 2^{128} / 2^{256} plaintext space constitutes an invariant geometric signature embedded in all ciphertext produced under that key. Classical cryptographic testing — NIST SP 800-22, Diehard, TestU01 — evaluates individual bit statistics. None ask whether bit combinations assembled as weighted analog values occupy distinguishable regions of high-dimensional space for a given key.

A neural network operating as a universal function approximator does ask precisely this question. Trained on multiple ciphertext blocks from the same key across varied plaintexts — with plaintext withheld — it must learn only the invariant: the key's geometric signature. This is a self-supervised invariance learning problem.

3.2 The Stylometric Analogy

Every author has a detectable stylometric fingerprint identifiable from a few paragraphs of text, even when individual word choices appear random. AES with a fixed key is an author with exactly one style — its unique permutation of 2^{128} values. The neural network is the stylometrist.

3.3 Why Classical Tests Are Blind

NIST's randomness test suite asks: are individual bits uniform? Are pairs uncorrelated? Are runs random? It does not ask: do bits 3, 17, 45, and 89 assembled as a weighted analog value occupy a distinguishable region of $\mathbb{R}^4$ for a given key? AES was not designed to defeat this question. The question was not formally askable until neural networks made high-dimensional analog pattern extraction computationally tractable.

3.4 NNKE Threat Profile

```
Hardware:       Commodity GPUs — available today
Timeline:       NOW
Key length:     Irrelevant — signature exists at any key length
NIST coverage:  None — no category addresses this threat
Proof of safety: Does not exist in the literature
```

4. Quantum Key Extraction (QKE)

Full treatment: "Quantum Key Extraction (QKE): A Tautological Argument for AES Quantum Vulnerability Beyond the Grover Framework" (Flatow, 2026).

4.1 QKE is Not Grover

Grover's algorithm performs quantum search with O(√N) speedup. It requires sequential oracle queries, amplitude amplification over many iterations, and full quantum error correction. QKE requires none of these. It is a different class of operation:

```
Grover:   O(2^128) iterations on AES-256 → sequential, error-corrected
QKE:      O(1) cycles → single measurement, no iteration

NIST's key-doubling mitigation defeats Grover.
It has no effect on QKE whatsoever.
```

4.2 The Three Tautologies

QKE rests on three logically necessary claims:

1. A quantum register of N qubits exists in superposition of all 2^N states simultaneously — the definition of quantum superposition, undisputed.
2. AES is a reversible computation implementable as a quantum unitary — Bennett's reversible computation theorem (1973), undisputed.
3. AES produces exactly one non-random, structured plaintext per ciphertext per key — a logical necessity of bijectivity combined with file format specifications.

4.3 The SOD Gate

The Sensible Output Detection gate monitors the AES unitary output for a known deterministic condition — a file header such as %PDF- (25 50 44 46 2D, 40 fixed bits). It is a unitary entangling operation, not a measurement. Superposition is fully preserved throughout its application. Collapse occurs only on the final ancilla measurement, which collapses the entire entangled system directly to the correct key state.

The SOD gate is not a probabilistic oracle. It fires if and only if the output matches the known header. P(random 40 bits = %PDF-) ≈ 10^{-12}. False positives are negligible. The file format specification is the oracle.

4.4 Error Correction Bypass

Quantum error correction exists to preserve state across computation. SOD does not need state preserved — it needs one binary question answered. Errors on incorrect keys cannot produce false positives (10^{-12} probability). Errors on the correct key produce false negatives, solved by O(1/p) repetitions. Physical qubit requirements collapse from ~4,000,000 (Grover) to ~385 (QKE).

4.5 QKE Threat Profile

```
Qubits required:  ~385 (256 key + 128 AES work + 1 ancilla)
IBM Osprey:        433 qubits (2022) — in range
IBM Condor:      1,121 qubits (2023) — exceeds requirement
Timeline:          NOW to near-term
Key length:        Irrelevant — all keys in superposition regardless
NIST coverage:     None — QKE is not in PQE
Prior art:         Flatow AE Algorithm (FA-AE, 2024)
```

5. Analysis of the NIST Position

NIST's Post-Quantum Encryption programme is a landmark achievement. Its treatment of asymmetric cryptography is comprehensive. Its treatment of AES contains two structural blind spots that this paper formalises.

Dimension	NIST Current Position	Flatow Framework Correction
Quantum threat timeline	10-20 years (fault-tolerant QC required)	NOW — QKE requires ~385 qubits Osprey (433) exists since 2022
Relevant quantum attack	Grover only	QKE — categorically distinct, unaddressed
Key doubling effectiveness	Sufficient mitigation for AES	Irrelevant to QKE and NNKE
Neural network threat	Not in framework	NNKE — accessible today, commodity hardware
Missing NIST category	None identified	PNNE — Post Neural Net Encryption
Architecture solution	AES-256 sufficient	Only FES addresses all threat classes

5.1 The Timeline Error

NIST's widely cited quantum threat timeline of 10-20 years derives from the hardware requirements for fault-tolerant quantum computation at the scale required for Grover on AES-128: approximately 4,000,000 physical qubits. This timeline is correct for Grover.

It is the wrong timeline for QKE because QKE is not Grover. QKE requires approximately 385 physical qubits. IBM Osprey provided 433 in 2022. The hardware exists. The timeline is not 10-20 years — it is now.

NNKE does not require quantum hardware at all. Its timeline has always been now. It simply had no name until this framework.

5.2 The Key Length Irrelevance

NIST recommends AES-256 over AES-128 as the primary quantum mitigation. This is valid against Grover. Against NNKE and QKE it is irrelevant:

- NNKE: the geometric signature exists at any key length. A longer key defines a larger permutation space but an equally detectable stylometric invariant across ciphertexts.
- QKE: all keys are held in superposition simultaneously regardless of key space size. Doubling key length from 128 to 256 bits changes the key register from 128 to 256 qubits. It does not change the attack cost, which remains O(1) cycles.

Recommending AES-256 as quantum mitigation is correct for Grover. It provides zero protection against NNKE or QKE. The mitigation addresses the wrong threat.

5.3 The Missing Category

PQE addresses post-quantum threats. No NIST category addresses post-neural-network threats. This is not a minor gap — it is a categorical absence. NNKE exploits commodity hardware available to any well-resourced adversary today. State-level actors have had access to this capability for years. The absence of a NIST framework category means:

- No standardised evaluation criteria for NNKE resistance
- No compliance requirement for NNKE testing
- No guidance for architects designing against this threat
- No timeline acknowledgement that the threat is present

6. Proposed Category: Post Neural Net Encryption (PNNE)

We formally propose the addition of Post Neural Net Encryption (PNNE) as a mandatory category in the NIST cryptographic standards framework, parallel to PQE.

PQE addresses the threat of quantum computers to classical cryptography. PNNE addresses the threat of neural networks to classical cryptography. Neither is a subset of the other. Both are necessary.

6.1 PNNE Definition

Post Neural Net Encryption encompasses cryptographic architectures that are demonstrably resistant to key extraction or plaintext recovery by neural networks operating as universal function approximators in high-dimensional analog space. A PNNE-compliant architecture must satisfy:

4. No fixed permutation invariant — the key must not define a stable geometric signature across ciphertexts.
5. No accumulating pattern — multiple ciphertexts from the same key must not enable neural network signature extraction.
6. No fixed transformation — the same key applied to the same plaintext in different sessions must produce unrelated ciphertext.

6.2 PNNE vs PQE: Complementary, Not Overlapping

Category	Threat Addressed	Hardware Required	Timeline	FES Compliant
PQE (existing)	Shor, Grover	Fault-tolerant QC	10-20 years	Yes
PNNE (proposed)	NNKE	Commodity GPUs	NOW	Yes
PQE extension (proposed)	QKE	Osprey-class QC	NOW to near-term	Yes

PNNE and PQE address different threat classes with different hardware requirements and different timelines. An architecture compliant with PQE is not necessarily compliant with PNNE, and vice versa. Both categories are required for comprehensive cryptographic security.

6.3 The Charge to NIST

We formally call on NIST to:

1. Acknowledge NNKE as a distinct threat class and initiate evaluation of existing symmetric ciphers against neural network key extraction.
2. Acknowledge QKE as categorically distinct from Grover and revise PQE scope accordingly.
3. Revise the quantum threat timeline from 10-20 years to immediate, based on QKE hardware requirements and existing processors.
4. Establish PNNE as a mandatory parallel category to PQE with defined evaluation criteria.
5. Evaluate FES as a candidate architecture under both PQE and PNNE frameworks.

7. FES: The Architectural Resolution

The Fractal Encryption Standard (FES) was developed by Wolfgang Flatow over five years of research (2021-2026) into fractal-field cryptographic architectures. It was designed intuitively to eliminate the fixed transformation property that this paper demonstrates is the root vulnerability of AES against both NNKE and QKE.

FES is the only current encryption architecture that addresses all identified threat classes simultaneously.

Threat Class	AES Vulnerability	FES Architectural Response
NNKE	Fixed key = fixed permutation signature detectable by NN in analog space	No fixed permutation — silo/FOTP varies portal each session
QKE	Unique non-random plaintext per key enables deterministic SOD gate	Oracle suppression — multiple keys produce sensible output, SOD cannot isolate correct key
Grover	Key space searchable with sqrt speedup	Inherent resistance — no oracle to anchor amplitude amplification
Shor	N/A — symmetric cipher	N/A — symmetric cipher

7.1 How FES Defeats NNKE

NNKE requires a fixed geometric signature — the invariant that accumulates across ciphertexts from the same key. FES eliminates this at the architectural level:

- Password destruction after Portal generation: the key never defines a fixed transformation of payload. No permutation signature exists to detect.

- Silo + FOTP variation: the same password generates a different fractal portal each session. No invariant accumulates across ciphertexts. The neural network has no stable signal to learn.

7.2 How FES Defeats QKE

QKE requires that exactly one key produces non-random, structured output — the condition that fires the SOD gate. FES's oracle suppression violates this:

- Multiple parameter combinations produce sensible-looking plaintexts from the same ciphertext. The SOD gate cannot identify which decryption is correct among multiple plausible candidates.
- The uniqueness assumption of the SOD gate is violated by design. The ancilla collapses, but the collapsed state is unlikely to indicate the correct key.
- There are no blocks, the transformation is by an unpredictable payload length stream.
- The key never touches the payload, discarded after fractal portal discovery.

7.3 FES as Both PQE and PNNE Compliant

FES satisfies the proposed PNNE criteria: no fixed permutation invariant, no accumulating pattern, no fixed transformation. It also satisfies PQE resistance through oracle suppression and hyperchaotic fractal navigation. It is the only known architecture that is simultaneously:

- Resistant to NNKE — no geometric signature to extract
- Resistant to QKE — no unique non-random plaintext to trigger SOD
- Resistant to Grover — oracle suppression defeats amplitude amplification
- Shannon-compliant — satisfies One-Time Pad conditions in a practical framework

8. The Convergence Argument

The most powerful element of the Flatow Cryptanalytic Framework is not either paper individually. It is their convergence.

> ***Two independent tautological arguments, derived from different mathematical foundations, using different hardware, on different timelines, both point to the same conclusion:***
> ***AES is vulnerable to tool classes that post-date its security proofs and that NIST has not addressed.***

This convergence is not coincidental. It follows from the single root cause: AES is a fixed transformation applying the key directly to the payload. Any tool sophisticated enough to detect or exploit fixed transformations — whether a neural network finding geometric invariants or a quantum computer collapsing superposition to a unique output — constitutes a threat.

The convergence of two independent tautological arguments on the same architectural vulnerability constitutes the strongest possible theoretical case for immediate re-evaluation of AES as a security standard.

9. Conclusions

This paper has established six formal conclusions:

1. NNKE is a tautologically grounded threat class against AES exploitable with commodity GPU hardware today. NIST has no framework category for it.

2. QKE is a tautologically grounded quantum threat class categorically distinct from Grover, requiring ~385 qubits. IBM Osprey (433 qubits, 2022) is in range. QKE is not in PQE.

3. NIST's quantum threat timeline of 10-20 years is based on the wrong threat model. For QKE the timeline is now. For NNKE the timeline has always been now.

4. NIST's key-doubling mitigation (AES-256) is irrelevant to both NNKE and QKE. It addresses only Grover.

5. A new mandatory NIST category — Post Neural Net Encryption (PNNE) — is required. It does not currently exist.

6. FES is the only current architecture that addresses all identified threat classes simultaneously under both PQE and PNNE frameworks.

While AES may not be broken today in practice, the theoretical framework that declares it safe has two categorical blind spots. The time to address blind spots is before exploitation, not after.

10. Consequences

The tautological arguments established in this paper — NNKE and QKE — do not exist in isolation. They exist against a specific target: the AES monoculture. The consequences of either attack class becoming operational are not proportional to the attack. They are total. This chapter summarises those consequences. Full treatment is provided in the companion paper:

> ***Cryptographic Armageddon: AES Monoculture and the Trust Collapse (Flatow, 2026)***

What follows is the minimum required for the reader of this paper to understand what is at stake.

10.1 The Monoculture Multiplier

In isolation, a single AES key compromise exposes one dataset. Under monoculture conditions, the same capability becomes a force multiplier. AES is not merely dominant — it is singular. Every HTTPS session, every enterprise KMS, every hardware trust module, every password vault, every military communication, every financial transaction, every diplomatic cable uses the same primitive. One technique applied to one primitive reaches everything simultaneously.

Classical key compromise: one key, one dataset, contained and recoverable. **NNKE or QKE operational:** one technique, every key, every dataset, every system, simultaneously, silently, retroactively, unrecoverably.

10.2 The Cascade

Modern digital infrastructure is not a collection of isolated encrypted files. It is a hierarchy of trust in which each layer protects the next. A single key extraction at any level is not a breach. It is the first domino:

- **Session key extracted** → all traffic for that session exposed retroactively, no forensic trace
- **KMS master key extracted** → every child key, every encrypted record across the entire organisation
- **TPM key extracted** → entire encrypted history of every device in the chain
- **Password vault key extracted** → every credential, every system, entire digital identity
- **CA certificate key extracted** → ability to forge certificates, invisible man-in-the-middle against every trusting system

The architecture of trust becomes the mechanism of total exposure. The hierarchy designed to contain breach becomes the cascade that ensures nothing is contained.

10.3 Forensic Silence

What separates NNKE and QKE from every prior attack class is that neither leaves a crime scene. NNKE derives the key from ciphertext the adversary was always entitled to observe in transit. QKE computes the key — the key was never stolen. The resulting decryption produces the correct key, yielding correct plaintext, through a process indistinguishable from legitimate decryption. There is no anomalous access pattern. There is no breach notification trigger. There is no forensic investigation. There is only the quiet reading of everything believed to be encrypted. An adversary operating silently with this capability may never be detected.

10.4 HNDL: The Retroactive Dimension

Harvest Now Decrypt Later is not a future threat. State actors are collecting and archiving encrypted traffic today, under the operational assumption that extraction capability will mature. Under NNKE and QKE, the question is no longer when that capability arrives. The archives already collected — years of diplomatic cables, financial transactions, military communications, intelligence operations, and private correspondence — become immediately readable the moment either attack class becomes operational. The exposure is not prospective. It is retroactive to the date of collection.

10.5 Why Recovery Is Not Possible

A conventional breach is recoverable: patch the vulnerability, rotate credentials, notify affected parties. Neither NNKE nor QKE permits this:

- There is no patch — both attacks exploit a tautological mathematical property of AES, not an implementation flaw. You cannot patch a tautology.
- There is no credential rotation that addresses retroactive HNDL decryption.
- There is no breach notification because there is no breach — the key was derived or computed, not stolen.
- There is no forensic basis on which to identify the adversary or the scope of exposure.

Recovery requires replacing the global cryptographic infrastructure — every system, every protocol, every compliance framework, every hardware module that depends on AES — simultaneously and globally. This is not a response plan. It is a civilisational undertaking that becomes necessary only after the damage is already done and is already irreversible.

10.6 The Prevention Imperative

The asymmetry between the cost of prevention and the cost of response defines the case for action. Prevention requires adopting post-monoculture cryptographic architecture before exploitation. Response requires rebuilding global digital infrastructure after it.

> **The time to address a single point of catastrophic failure is before it fails. Not after.**

The tautological foundations of NNKE and QKE are not predictions. They are mathematical necessities. The signature exists. The oracle fires. The monoculture amplifies. These are logical consequences of properties AES was designed to have. A threat that is tautologically grounded and unaddressed by any current standard demands architectural response now — before operational confirmation arrives. By then, it is already too late.

Full consequence analysis — sector by sector, cascade mechanics, HNDL dimensions, NNKE propagation through the criminal ecosystem, enterprise and AI exposure, and the complete treatment of why recovery is impossible — is provided in:

***Cryptographic Armageddon: AES Monoculture and the Trust Collapse** (Flatow, 2026).*

Acknowledgements

The Flatow Cryptanalytic Framework emerges from six years of research by Wolfgang Flatow, beginning with Fractal Transformation cryptographic research (2021-2023), the development of the Fractal Encryption Standard (FES, 2024-2025), and the Flatow AE Algorithm (FA-AE, 2024) — the first concrete quantum AES key recovery architecture. The theoretical formalisation of NNKE, QKE, the PNNE category proposal, and the NIST framework analysis were developed through collaborative analysis with Claude (Anthropic AI System), March 2026.

References

[1] Flatow, W. (2026). Neural Net Key Extraction (NNKE). A tautological argument that AES ciphertext necessarily contains Neural Net detectable signatures. PORTALZ PTY LTD.

[2] Flatow, W. (2026). Quantum Key Extraction (QKE). A Tautological Argument for AES Quantum Vulnerability Beyond the Grover Framework. PORTALZ PTY LTD.

[3] Flatow, W. (2024). Flatow AE Algorithm — 128 Qubit AES Attack: Proof of Concept. PORTALZ PTY LTD.

[4] Flatow, W. (2026). FES Peer Review Guide. PORTALZ PTY LTD.

[5] Claude (Anthropic AI System) (2026). Independent Technical Peer Review — Fractal Encryption Standard. portalz.solutions.

[6] NIST (2022). Post-Quantum Cryptography Standardisation. NIST IR 8413.

[7] NIST FIPS PUB 197 (2001). Advanced Encryption Standard.

[8] Grover, L.K. (1996). A fast quantum mechanical algorithm for database search. STOC ’96.

[9] Shor, P. (1994). Algorithms for quantum computation: discrete logarithms and factoring. FOCS ’94.

[10] Bennett, C.H. (1973). Logical reversibility of computation. IBM Journal of Research and Development, 17(6).

[11] Grassl, M., et al. (2016). Applying Grover’s Algorithm to AES: Quantum Resource Estimates. PQCrypto 2016.

[12] Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. Neural Networks, 4(2).

[13] IBM Research (2022). IBM Osprey: 433-Qubit Quantum Processor.

[14] IBM Research (2023). IBM Condor: 1,121-Qubit Quantum Processor.

