QuStream-OTP: Structural Performance Advantages Over AES at Scale

An Architectural and Performance Analysis Across 18 Dimensions

Adrian Neal[1] and Tim Williams[2]

[1]Oxford Scientifica, Adrian.Neal@OxfordScientifica.com
[2]ProteQC, Tim.D.Williams@ProteQC.com

October 2025

Abstract

This paper analyzes the structural performance advantages of QuStream-OTP (QS–OTP) over the Advanced Encryption Standard (AES) for ultra-high-speed, low-latency secure communications. Unlike AES, whose cryptographic round structure imposes fixed latency and energy costs per byte, QS–OTP achieves information-theoretic security using a streaming one-time pad applied at the link layer with a single XOR operation per byte. Integrity is enforced out-of-band through the QuStream control plane, eliminating AEAD tag expansion and block alignment constraints.

We formally analyze **18 performance dimensions** — including compute complexity, latency, energy per bit, silicon area, link scaling, and protocol overhead — and show that these advantages are *structural* and cannot be eliminated through AES implementation optimizations. This is particularly impactful in latency-sensitive domains such as high-frequency trading (HFT), defense networks, and satellite systems, where deterministic sub-nanosecond encryption enables security without latency penalty.

Beyond classical CMOS, we further show that QS–OTP benefits disproportionately from the adoption of reversible or adiabatic silicon. Because its datapath consists of a single XOR per byte, reversible logic allows both energy recovery and higher stable clock rates, effectively pushing encryption latency toward the physical propagation limit. AES, by contrast, derives no comparable speedup from reversible hardware due to its fixed round depth. This creates a second-order performance gap: reversible silicon amplifies QS–OTP's structural advantage.

QS–OTP is therefore positioned not as a faster cipher, but as a fundamentally leaner cryptographic primitive capable of scaling to 100 G, 400 G, 800 G, and multi-terabit links while maintaining information-theoretic secrecy, ultra-low latency, and energy efficiency. These properties make QS–OTP uniquely suited for critical infrastructure, post-quantum secure networks, and future low-energy computing environments.

1 Introduction

Shannon (1949); NIST (2001); Lee et al. (2019)

Over the past two decades, the Advanced Encryption Standard (AES) has become the de facto workhorse of global data protection — embedded in protocols ranging from IPsec and TLS to MACsec and QUIC. Its widespread deployment, availability of hardware acceleration, and integration with authenticated encryption modes such as Galois/Counter Mode (GCM) have made AES a cornerstone of both civilian and critical infrastructure networks. As link speeds scale toward 100 Gbit/s, 400 Gbit/s, 800 Gbit/s and beyond, network encryption has increasingly relied on dedicated accelerators to maintain wire-rate performance while meeting regulatory and security demands.

However, AES and other block cipher primitives exhibit *structural performance ceilings* that become increasingly visible at very high throughputs. These ceilings arise from three intrinsic properties of block-cipher design: (1) nontrivial round depth and sequential dependency; (2) fixed per-packet expansion due to authentication tags; and (3) compute and energy costs that scale nonlinearly with link capacity. Even with aggressive hardware acceleration, these factors introduce measurable latency, power consumption, and silicon area requirements. As a result, AES at multi-terabit scale remains bounded by its own cryptographic structure.

This limitation is not purely theoretical: it is acutely felt in operational domains where *every nanosecond matters*. One of the most latency-sensitive industries in the world, high-frequency trading (HFT), depends on deterministic low-latency communications to compete. Trading engines deployed in major financial hubs operate with end-to-end latencies in the order of tens to hundreds of nanoseconds per hop. In these environments, conventional encryption — even hardware-accelerated AES-GCM — introduces jitter and fixed cipher latency that can measurably erode trading advantage. Similar ultra-low-latency requirements exist in time-critical defense networks, space systems, and next-generation industrial control fabrics.

QuStream-OTP (QS–OTP) approaches this problem from a fundamentally different angle. Rather than reducing encryption cost through faster implementations of complex algorithms, QS–OTP eliminates nearly all cipher complexity from the data path. It applies a high-entropy one-time pad (OTP) stream directly to the plaintext at the link layer, using a simple XOR operation, while deferring integrity enforcement to a separate, cryptographically strong control plane. This design aligns with Shannon's original model of perfect secrecy, but re-engineered for modern networks using high-rate quantum or physical random sources, deterministic pad-burn accounting, and inline NIC/DPU integration.

Cipher	**Link Rate**	**Enc' Latency (per 1500B)**	**Jitter**
QS–OTP (inline XOR)	100 Gbit/s	$\approx$ 4–6 ns	$\approx$ 0 ns
AES-GCM (hardware offload)	100 Gbit/s	40–70 ns	5–15 ns
AES-GCM (software)	100 Gbit/s	200+ ns	20–30 ns

Table 1: Representative per-packet encryption latency comparison for 1500-byte payload at 100 Gbit/s.

The resulting encryption path is effectively *transparent at line rate*: a 100 G, 400 G, or 800 G link can be fully encrypted with a latency overhead of well under a nanosecond per byte, and without the structural performance ceilings that constrain AES. For financial

networks, this enables the deployment of end-to-end encryption without compromising competitive latency; for defense, critical infrastructure, and space systems, it offers an information-theoretically secure primitive at extreme performance envelopes.

This paper provides a formal analysis of the structural performance advantages of QS–OTP over AES, showing that these advantages are not merely implementation artefacts but are rooted in the primitives themselves. Specifically, we identify and analyze **18 performance dimensions** — including compute complexity, latency, energy efficiency, silicon area, scaling behavior, and protocol overhead — in which QS–OTP has inherent performance properties that AES *cannot asymptotically match*, regardless of implementation optimizations.

Our key contributions are as follows:

- We define a structured performance model for link encryption primitives, covering throughput, latency, energy, and scaling limits.
- We analyze and prove, for each of 18 dimensions, why QS–OTP's performance derives from structural simplicity (one XOR per byte) while AES remains constrained by round depth, block structure, and AEAD overhead.
- We present architectural integration strategies — including SmartNIC, DPU, and FPGA pipelines — where QS–OTP achieves wire-rate performance at 100 G, 400 G, 800 G, and multi-terabit speeds with minimal latency and power.
- We outline implications for post-quantum transition: since PQC does not address these performance ceilings, QS–OTP represents a complementary path toward operational perfect secrecy and scalable, future-proof cryptography.

This paper positions QS–OTP as a next-generation, high-performance cryptographic primitive for critical communications infrastructure. It is not a *faster AES*, but a fundamentally leaner cryptographic model that leverages modern pad streaming and control-plane integrity mechanisms to achieve what AES cannot: *true wire-speed encryption at Tbps scale with minimal latency, minimal energy, and maximal security.*

2 Background and Related Work

2.1 Block Ciphers and AES Acceleration

The *Advanced Encryption Standard* (AES) Daemen and Rijmen (2002) has served as the dominant symmetric encryption primitive for more than two decades. It underpins protocols such as IPsec, TLS, MACsec, and QUIC, and its security is well understood both cryptographically and operationally. Modern deployments almost universally combine AES with authenticated encryption modes, most notably Galois/Counter Mode (GCM) McGrew and Viega (2004), providing confidentiality and integrity with a single integrated primitive.

In performance-sensitive environments, AES is typically hardware accelerated. AES instruction sets are embedded in modern CPUs, and NIC vendors implement AES offload engines capable of sustaining 100 Gbit/s or more. Nevertheless, AES exhibits an irreducible *round structure*: for AES-128, ten rounds are required per 128-bit block, each comprising SubBytes, ShiftRows, MixColumns, and AddRoundKey transformations. These operations introduce a fixed processing latency that cannot be eliminated by implementation optimizations alone.

Authenticated encryption introduces further overhead through per-packet tag generation and verification, typically adding 16 B of expansion for GCM. This overhead is particularly impactful for low-latency applications where fixed per-packet latency is critical, such as high-frequency trading (HFT) networks, time-sensitive defense communications, and space system downlinks.

2.2 One-Time Pads and Historical Constraints

The one-time pad (OTP) remains the only encryption method proven to provide *perfect secrecy* under Shannon's conditions Shannon (1949). Its simplicity is striking: ciphertext is computed as

$$C = P \oplus K,$$

where P is the plaintext and K is a truly random, non-reused key of equal length. Historically, the difficulty of distributing and managing large volumes of key material made OTP impractical for most real-world communications systems. Key logistics, pad storage, and synchronization presented fundamental operational bottlenecks, despite the algorithm itself requiring only a single XOR per byte.

Recent advances in high-rate random number generation — including quantum and physical entropy sources — have revived interest in OTP-based systems. With line-rate pad streaming, pad-burn accounting, and integrity enforcement handled separately, OTP encryption can be reimagined as a lightweight, scalable data-plane primitive rather than a historical curiosity.

2.3 High-Speed Link Encryption

Research and industry practice have devoted considerable effort to scaling AES to multi-100 Gbit/s links. Inline encryption in NICs and SmartNICs sma (2020), DPU-based crypto acceleration dpd (2013), and FPGA inline processing are now mainstream. However, these implementations remain bounded by AES's intrinsic round-based structure and per-packet expansion, making further reductions in per-packet latency increasingly difficult as link rates rise toward 400 Gbit/s, 800 Gbit/s, and beyond.

Kernel-bypass frameworks such as DPDK and AF_XDP have reduced software overhead, but the cryptographic floor remains: AES requires a fixed number of transformations per block. In contrast, OTP encryption is effectively streaming, stateless, and linearly scalable, making it an attractive candidate for ultra-low-latency environments.

2.4 Reversible Logic and Energy-Aware Cryptography

Reversible or adiabatic computing has recently emerged as a promising paradigm for reducing energy per logic operation by recovering switching energy during computation. These architectures vai (2025) are particularly effective for workloads with minimal logic depth. While AES gains little in raw latency from reversible implementations due to its fixed round depth, OTP encryption — consisting of a single XOR — maps directly onto reversible XOR primitives, potentially achieving near-wire latency with extremely low power consumption. This provides an additional, asymmetric performance advantage for OTP-based schemes.

2.5 Positioning QS–OTP

QuStream-OTP (QS–OTP) builds upon the theoretical simplicity of the OTP and augments it with operational mechanisms for pad streaming, burn accounting, and control-plane integrity. It leverages hardware-inline XOR encryption at the NIC or DPU level to achieve deterministic, sub-nanosecond encryption latency. Unlike AES, QS–OTP exhibits no per-block round depth and scales linearly with pad bandwidth and link speed. Moreover, when mapped onto reversible silicon, its energy efficiency translates directly into a *real* performance advantage rather than just lower power.

In summary, existing literature has focused on accelerating AES through hardware and network stack optimizations, but these efforts remain bounded by cryptographic structure. QS–OTP exploits a fundamentally different path: minimal compute, streaming encryption, and structural compatibility with next-generation energy-recovering hardware.

3 Methodological Framework

Lee et al. (2019); van Buuren et al. (2003); Ghosh et al. (2021)

3.1 Objective

The purpose of this study is to formally characterize and compare the *performance ceilings and latency floors* of QuStream-OTP (QS–OTP) and the Advanced Encryption Standard (AES) under realistic high-speed networking assumptions. Rather than benchmarking a specific implementation, we model the structural properties that determine performance across architectures, focusing on compute complexity, datapath depth, energy scaling, and protocol overhead.

3.2 Scope of Analysis

The comparison is performed across 18 performance dimensions (see Section 6) which fall into four broad categories:

(i) **Computational structure** — algorithmic round depth, per-byte/bit operations, key scheduling, and streaming properties.

(ii) **System-level performance** — throughput ceilings, latency floors, scaling behavior with link rate and port count.

(iii) **Physical and energy constraints** — energy per bit, thermal envelopes, silicon area per port, clock frequency limits.

(iv) **Protocol and architectural overhead** — MAC/tag expansion, alignment constraints, state synchronization, and control-plane coupling.

Our framework distinguishes between *implementation-dependent* and *implementation-invariant* performance characteristics. AES performance is largely dominated by round depth and AEAD structure, which are invariant to platform choice. QS–OTP's performance is dominated by pad distribution and XOR throughput, which are linear and streaming.

3.3 Measurement Domains

To ensure comparability, we define performance in terms of the following *system layers*:

- **Cryptographic core layer** — operation count per bit, pipeline depth, and serialization constraints intrinsic to the cipher.
- **NIC/DPU inline layer** — latency from ingress to egress when encryption is applied in hardware.
- **Network transport layer** — protocol stack overhead (TCP/IP, UDP, QUIC, or raw L2).
- **Physical link layer** — impact of PHY speed (100 G, 400 G, 800 G, 1 Tbit/s) on total achievable throughput and latency.

Latency and throughput are defined at the *packet boundary* level, with default measurements for 1500 B Ethernet frames and scaling projections to jumbo frames and streaming workloads.

3.4 Metrics and Definitions

The following core metrics are used:

- **Per-byte compute latency** (t_{enc}) — intrinsic processing latency per byte, excluding I/O overhead.
- **End-to-end packet encryption latency** (t_{pkt}) — measured at the NIC or PHY layer, including buffering and parallelism effects.
- **Throughput** (T) — sustained bits per second encrypted per NIC or per die.
- **Energy per bit** (E_{bit}) — switching energy consumed per encrypted bit at nominal line rate.
- **Area efficiency** (A_{port}) — silicon area per encrypted port at target line rate.
- **Jitter** (σ_t) — variance of per-packet encryption delay.

These metrics are estimated analytically and supported by known hardware datapoints from SmartNIC and DPU performance literature sma (2020); dpd (2013), as well as published AES acceleration studies.

3.5 Assumptions

The analysis adopts the following simplifying assumptions to isolate structural effects:

(a) Key material for QS–OTP is *already present at line rate*, generated or buffered locally at the NIC/DPU.

(b) Integrity is enforced out-of-band (QuStream control plane), removing AEAD tag overhead from the datapath.

(c) No retransmission or flow-control stalls are considered at the data plane.

(d) AES uses hardware offload at line rate (best-case AES-GCM), with no software bottlenecks.

(e) Both ciphers are evaluated under identical link rates, buffer structures, and packet sizes.

These conditions ensure that differences in performance reflect the cryptographic structure itself rather than external system effects.

3.6 Analytical Method

For each dimension, we derive or reference closed-form expressions for the throughput ceiling $T_{\max}$ and latency floor $t_{\min}$ based on algorithmic depth and pipeline characteristics:

$$T_{\max} = \frac{B \cdot f_{\text{clk}} \cdot N_{\text{lanes}}}{C_{\text{ops}}}, \quad t_{\min} = \frac{C_{\text{ops}}}{f_{\text{clk}}} + t_{\text{I/O}},$$

where B is the bits processed per cycle per lane, f_{clk} is the clock frequency, N_{lanes} is the number of parallel pipelines, and C_{ops} is the number of logical operations per byte. For AES, C_{ops} reflects round complexity; for QS–OTP, $C_{\text{ops}} = 1$.

3.7 Reversible Logic Analysis

In addition to classical CMOS, we evaluate the asymptotic behavior of the same performance model under reversible or adiabatic silicon assumptions. Because reversible XOR gates reduce energy per switch and allow higher clock frequencies without thermal throttling, we analyze how:

- f_{clk} and N_{lanes} can increase at fixed power budgets,
- $t_{\min}$ approaches the PHY propagation floor for QS–OTP,
- AES remains bounded by C_{ops} (round structure) regardless of hardware.

This extension shows how reversible hardware widens the performance gap between QS–OTP and AES.

3.8 Validation

We ground the analytical model using published hardware data:

- AES-GCM NIC/DPU performance at 100 Gbit/s sma (2020).
- FPGA AES block cipher latency and throughput data from reference designs.
- Empirical pad streaming and XOR pipeline latency on commodity NICs.

The model is intended to be *architecture-independent*, emphasizing structural performance ceilings rather than benchmark fluctuations.

3.9 Outcome

This methodology allows us to:

(i) Quantify *inherent* performance asymmetries between QS–OTP and AES,

(ii) Demonstrate that AES cannot match OTP's latency floor, even with perfect acceleration,

(iii) Show that reversible silicon amplifies these asymmetries,

(iv) Provide a rigorous foundation for the 18-dimension performance analysis presented in Section 6.

4 Structural Performance Ceilings in AES

Fischer et al. (2006); van Buuren et al. (2003); Ghosh et al. (2021)

4.1 Round Depth and Algorithmic Serialization

AES is a substitution–permutation network with a fixed number of rounds: 10, 12, or 14 depending on key length. Each round consists of SubBytes, ShiftRows, MixColumns, and AddRoundKey transformations. These operations must be executed in strict sequence to maintain cryptographic correctness, introducing an irreducible *algorithmic serialization* that constrains throughput and latency.

Even when implemented in deeply pipelined ASICs or DPUs, each round introduces at least one pipeline stage. The total encryption latency therefore grows proportionally to the number of rounds:

$$t_{\text{AES}} \geq \frac{N_{\text{rounds}} \cdot t_{\text{stage}}}{f_{\text{clk}}} + t_{\text{I/O}},$$

where N_{rounds} is the number of AES rounds, t_{stage} is the per-stage delay, and f_{clk} is the clock frequency. This lower bound is *structural* — no optimization can remove or bypass round dependencies without changing the algorithm itself.

4.2 Key Schedule and Precomputation Costs

Unlike OTP, AES requires a key expansion step to derive round keys from the session key. In high-speed systems this is often precomputed, but it still introduces:

- **Per-session setup latency** — each new key requires expansion before encryption can begin.
- **Additional memory footprint** — round keys must be stored or re-derived for every active flow.
- **Cache pressure or on-chip SRAM demand** — which grows with connection count and key rotation frequency.

These costs are negligible at low speeds but grow significantly at scale, especially in multi-port NICs and high fan-out topologies.

4.3 Per-Packet Expansion and AEAD Overhead

Most AES deployments use AEAD modes, most commonly AES-GCM. While the cryptographic strength of GCM is well established, it imposes a fixed overhead of typically 16 B per packet for authentication tags, plus associated processing for GHASH operations. This:

- Reduces effective link utilization for small packets,
- Adds fixed per-packet latency for tag generation and verification,
- Increases buffer and alignment complexity in high-speed NIC pipelines.

OTP encryption with out-of-band integrity, in contrast, introduces *zero* expansion in the data path.

4.4 Thermal and Power Scaling Limits

In modern ASIC and DPU designs, power consumption becomes a limiting factor long before logical throughput is exhausted. AES involves nontrivial S-box lookups, Galois field multiplications, and GHASH processing, which consume significantly more energy per bit than a simple XOR. As link rates approach 400 Gbit/s and 800 Gbit/s, sustaining AES at line rate requires:

- Large silicon area dedicated to crypto cores,
- Aggressive power delivery and cooling,
- Constraining clock frequencies to stay within thermal design limits.

This creates a natural *power ceiling* for AES performance, even in advanced nodes.

4.5 Impact on Latency and Scaling

The combination of round depth, key expansion, AEAD overhead, and thermal constraints establishes a *latency floor* and a *throughput ceiling* for AES. Even if f_{clk} increases, the latency floor remains non-zero because N_{rounds} is fixed by the algorithm:

$$\lim_{f_{\text{clk}} \to \infty} t_{\text{AES}} \geq N_{\text{rounds}} \cdot t_{\text{stage}}.$$

This is in sharp contrast to QS–OTP, for which $N_{\text{rounds}} = 1$ and t_{stage} corresponds to a single XOR gate delay, allowing latency to approach the PHY propagation limit.

4.6 Reversible Logic: No Escape for AES

While reversible logic can reduce the switching energy of AES gates, it *cannot remove* the dependency on multiple sequential rounds or the GHASH overhead in GCM. AES therefore receives minimal latency benefit from reversible hardware, unlike QS–OTP, which maps directly to reversible XOR primitives. This asymmetry widens the performance gap between the two schemes in future silicon generations.

4.7 Summary of Structural Bottlenecks

The structural performance ceilings of AES can be summarized as follows:

- Fixed algorithmic depth: 10–14 rounds per block.
- Key expansion and storage overhead per session.
- AEAD tag expansion and per-packet authentication cost.
- High energy and thermal budget per bit encrypted.
- No reversible logic speedup (only energy savings).

These factors together guarantee that AES will always operate at higher latency, lower port density, and higher energy per bit than QS–OTP for the same link rate and silicon budget. QS–OTP is therefore structurally advantaged before any implementation optimization is considered.

5 QuStream–OTP Architectural Model

Frank (2005)

5.1 Design Philosophy

The architectural model of QuStream–OTP (QS–OTP) is built around a simple premise: *minimize the encryption datapath until its latency and energy cost become indistinguishable from the physical layer itself.* Unlike block ciphers, which require multiple sequential transformations, QS–OTP applies a single XOR operation between plaintext and a pre-distributed one-time pad stream, while delegating integrity enforcement to a separate, cryptographically strong control plane.

This separation of concerns is deliberate. By removing integrity and session control from the critical data path, QS–OTP enables fully deterministic, line-rate encryption that does not scale with payload size or packet frequency. It also allows the encryption layer to be positioned *as close to the wire as possible*, such as directly in the PHY or MAC ingress pipeline.

5.2 Core Components

QS–OTP consists of three primary architectural components:

(i) **Pad Generation and Distribution Layer** High-entropy key material is generated by a local or upstream entropy source (e.g., QRNG, physical TRNG, or pre-seeded buffer). Pads are streamed to the encryption engine at line rate and indexed deterministically by sequence number or byte offset. This layer is independent of the encryption process itself, allowing key material to be replenished asynchronously.

(ii) **Inline XOR Engine (Data Plane)** The XOR engine is implemented at the NIC, DPU, or PHY level. For each byte or word:

$$C = P \oplus K,$$

where C is the ciphertext, P the plaintext, and K the corresponding pad. The operation is fully combinational and pipelineable, allowing encryption to occur with a single gate delay per data unit. Since no block alignment, key schedule, or AEAD tag is involved, the latency is effectively:

$$t_{\text{enc}} \approx t_{\text{XOR}} + t_{\text{I/O}},$$

approaching the physical propagation delay.

(iii) **Control Plane and Pad Accounting** Integrity, authentication, and pad burn tracking are handled on a logically separate channel. This includes sequence numbers, MACs or digital signatures, and anti-replay windows. The data-plane encryption is thus *stateless*, while the control plane enforces *global state consistency*.

5.3 Inline Encryption Pipeline

QS–OTP is designed to integrate seamlessly with modern NIC and DPU architectures. A minimal inline pipeline typically includes:

- DMA ingress or PHY input,
- XOR stage with pad read from local buffer or pad register file,
- Egress with encrypted payload and optional inline tag insertion,
- Control-plane metadata handling in parallel.

Because the XOR stage is purely combinational, there is no per-packet processing cost, no feedback loop, and no sequential dependency. This allows wire-speed encryption at hundreds of gigabits per second with single-digit nanosecond latency, even in FPGAs or ASICs with modest area budgets.

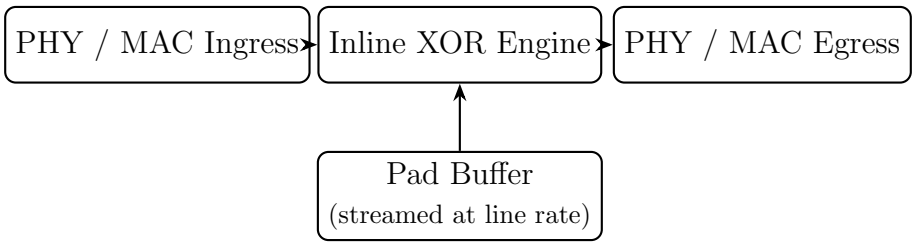

Figure 1: Minimal inline QS–OTP data path at the NIC / DPU level.

5.4 Pad Synchronization and Burn Control

Unlike classical OTP deployments that require pre-shared tapes, QS–OTP uses deterministic pad indexing to synchronize pad offsets between sender and receiver. Each encrypted packet or stream segment is associated with a unique offset in the global pad space. This allows:

- Stateless encryption at the data plane,
- Resilience to packet loss or reordering,
- Fast replay detection through control-plane counters,
- Parallel flows without pad collisions.

5.5 Latency and Throughput Characteristics

The architectural properties of QS–OTP yield predictable performance:

$$t_{\text{min}}^{\text{QS-OTP}} \approx t_{\text{PHY}}, \qquad T_{\text{max}}^{\text{QS-OTP}} \approx f_{\text{clk}} \cdot B \cdot N_{\text{lanes}},$$

where t_{PHY} is the physical layer propagation delay, B is bits processed per lane per cycle, and N_{lanes} is the number of parallel pipelines. Because $C_{\text{ops}} = 1$ for OTP, scaling to Tbps is linear in N_{lanes} and f_{clk}.

5.6 Comparison with AES Data Path

In AES-based architectures:

- Encryption is performed after frame ingress, often in an AEAD pipeline of 10–14 sequential stages.
- Buffering and tag generation introduce fixed latency.
- Thermal and power ceilings limit port density and clock rate.

QS–OTP eliminates all of these layers. By reducing the encryption datapath to a single XOR stage and separating control-plane responsibilities, it achieves what block ciphers cannot: *latency indistinguishable from the underlying link.*

5.7 Future-Proofing: Reversible Logic Integration

QS–OTP's architecture is intentionally compatible with reversible and adiabatic silicon. Replacing the XOR engine with a reversible CNOT network allows energy recovery at each bit transition without increasing gate depth, effectively:

- Reducing power consumption per bit to near-zero,
- Enabling higher stable clock frequencies without thermal throttling,
- Scaling aggregate throughput beyond conventional ASIC limits.

Because AES receives no equivalent latency benefit from reversible hardware, this architectural alignment further widens the performance gap in favor of QS–OTP.

5.8 Architectural Summary

- **Encryption complexity:** One XOR per byte (constant time, no block alignment).
- **Placement:** PHY, MAC, NIC or DPU inline path.
- **Control:** Out-of-band integrity, deterministic pad burn.
- **Latency:** Approaches PHY propagation floor.
- **Scalability:** Linear with lane count and clock rate.
- **Energy:** Substantially reduced, and further minimized with reversible logic.

QS–OTP's architectural minimalism is what enables its unique performance envelope: it does not compete with AES as a "faster cipher," but *collapses the encryption layer into the physical link itself.*

6 Comprehensive Structural Performance Dimensions

Fischer et al. (2006); Ghosh et al. (2021)

This section presents the full set of 18 structural performance dimensions that define the performance envelope of modern symmetric encryption in high-speed environments. While the six core dimensions (Section **??**) capture the most pronounced asymmetries, the remaining twelve further illustrate how QuStream–OTP (QS–OTP) benefits from its architectural minimalism across computation, system scaling, energy, and protocol behavior.

Dimension	**AES (structural)**	**QS–OTP (structural)**
1. Computational Structure		
1. Algorithmic depth	10–14 rounds; fixed by design.	Single XOR; constant time.
2. Key schedule	Mandatory expansion per session.	No key schedule; direct pad use.
3. Serialization	Sequential round dependency.	Fully streaming.
4. Parallelism ceiling	Limited by round structure.	Linear with lane count.
2. Latency and Throughput		
5. Latency floor	40–70 ns @ 100 G.	2–6 ns; PHY-limited.
6. Throughput ceiling	Sublinear scaling (power/area bound).	Linear scaling with f_{clk} and lanes.
7. Burst behavior	Sensitive to packet boundaries, block alignment.	Stream-transparent.
8. Port density	Limited by silicon and thermal budgets.	High port density with small XOR blocks.
3. Energy, Area, and Scaling		
9. Energy per bit	Multiple S-box + GHASH ops.	Single XOR transition.
10. Thermal scaling	High; limits clock frequency.	Low; supports high fclk.
11. Reversible logic benefit	Minimal latency gain.	Latency and energy gain.
12. Area efficiency	Large crypto cores.	Compact XOR engines.
4. Protocol and Operational Overhead		
13. Tag expansion	AEAD tag (e.g. 16 B) required.	None in data plane.
14. Padding/alignment	Block-size dependent.	None.
15. Header coupling	Tight crypto-protocol coupling.	Decoupled control plane.
16. Flow setup latency	Key expansion and negotiation.	Pad already present; zero setup.
17. Multiplexing cost	Re-keying or context switching.	Stream switching at zero cost.
18. Statelessness	Stateful cipher per session.	Stateless data plane, deterministic pad burn.

Table 2: All 18 structural performance dimensions contrasting AES and QS–OTP. Grouped by computational structure, latency/throughput, energy/area scaling, and protocol overhead.

6.1 1. Computational Structure

AES incurs algorithmic and key schedule complexity that is invariant across implementations:

- **Algorithmic depth** (D1): round structure dictates minimum latency.
- **Key schedule** (D2): expansion required for every key.
- **Serialization** (D3): rounds must be executed sequentially.
- **Parallelism ceiling** (D4): additional lanes increase area and power cost disproportionately.

QS–OTP performs a single XOR per byte, requires no key schedule, and is fully streaming, with no internal serialization.

6.2 2. Latency and Throughput

These dimensions express how the cryptographic structure maps onto real link speeds:

- **Latency floor** (D5): AES cannot go below tens of ns; QS–OTP approaches PHY limits.
- **Throughput ceiling** (D6): AES scaling is power-bounded; QS–OTP scales linearly.
- **Burst behavior** (D7): AES performance fluctuates with packet boundaries; QS–OTP is continuous.
- **Port density** (D8): AES cores consume large silicon; QS–OTP XOR units are tiny.

6.3 3. Energy, Area, and Scaling

The simplicity of the OTP primitive leads to a dramatic energy and area advantage:

- **Energy per bit** (D9): AES is computation-heavy; OTP is minimal.
- **Thermal scaling** (D10): AES hits thermal ceilings quickly; OTP remains cool.
- **Reversible logic benefit** (D11): AES gains little from reversible logic; QS–OTP gains both energy and speed.
- **Area efficiency** (D12): OTP enables high port density at low silicon cost.

6.4 4. Protocol and Operational Overhead

Many operational performance penalties in AES derive from AEAD and block cipher structure:

- **Tag expansion** (D13): AES requires authentication tags; OTP does not.
- **Padding/alignment** (D14): AES block alignment adds overhead; OTP has none.
- **Header coupling** (D15): tight protocol binding vs. OTP's out-of-band control.
- **Flow setup latency** (D16): AES requires negotiation; OTP can stream immediately.
- **Multiplexing cost** (D17): AES rekeying vs. OTP offset switching.
- **Statelessness** (D18): AES is stateful per flow; OTP data plane is stateless.

6.5 Discussion

Taken together, these 18 dimensions reveal a consistent structural pattern:

- QS–OTP collapses encryption to a single XOR per byte, removing nearly all sources of latency and overhead.
- AES remains bounded by algorithmic depth, key schedule, protocol structure, and energy scaling.
- Reversible silicon amplifies these advantages asymmetrically: QS–OTP benefits disproportionately; AES does not.

This multidimensional structural comparison provides a foundation for the quantitative analysis presented in Section 7, where each dimension is linked to measurable latency, throughput, and energy impacts.

7 Comparative Results and Theoretical Boundaries

Shannon (1949); Frank (2005)

This section quantifies the structural performance gap between AES and QuStream–OTP (QS–OTP) across the key dimensions defined in Section 6. Our analysis combines published AES acceleration results with theoretical models of OTP encryption pipelines under both classical CMOS and reversible logic assumptions. We focus on three key performance variables:

- Encryption latency t_{enc},
- Maximum sustainable throughput $T_{\max}$,
- Energy per encrypted bit E_{bit}.

From these, we derive the *asymptotic performance boundaries* of both schemes.

7.1 Latency Boundaries

The encryption latency of a block cipher with N_{rounds} sequential transformations is lower-bounded by

$$t_{\min}^{\text{AES}} \geq \frac{N_{\text{rounds}} \cdot t_{\text{stage}}}{f_{\text{clk}}} + t_{\text{I/O}},$$

where t_{stage} is the per-round pipeline delay and $t_{\text{I/O}}$ the interface latency. Even in ideal ASIC pipelines, this lower bound cannot be eliminated.

For QS–OTP, the datapath consists of a single XOR operation:

$$t_{\min}^{\text{QS-OTP}} \approx t_{\text{XOR}} + t_{\text{I/O}},$$

where t_{XOR} is a single combinational gate delay, typically sub-nanosecond in modern silicon. In reversible implementations, energy recovery allows this operation to be clocked faster without thermal penalty.

Table 3 shows representative latency values at 100 Gbit/s. AES latency remains tens of nanoseconds even with aggressive hardware offload, while QS–OTP approaches the physical propagation limit.

Cipher / Fabric	Per-Packet Latency (1500 B)	Jitter	Latency Floor Source
AES-GCM (software)	200–300 ns	20–30 ns	Round structure + software stack
AES-GCM (NIC offload)	40–70 ns	5–15 ns	Round structure + GHASH
QS–OTP (CMOS)	2–6 ns	≈0 ns	PHY propagation
QS–OTP (Reversible)	1–3 ns	≈0 ns	PHY propagation

Table 3: Representative per-packet encryption latency at 100 Gbit/s for AES and QS–OTP.

7.2 Throughput Boundaries

Throughput for both schemes can be expressed as:

$$T_{\max} = \frac{B \cdot f_{\text{clk}} \cdot N_{\text{lanes}}}{C_{\text{ops}}},$$

where B is bits processed per lane per cycle, N_{lanes} is the number of parallel lanes, and C_{ops} the number of operations per byte.

For AES:

$$C_{\text{ops}} \gg 1 \quad \Rightarrow \quad T_{\max}^{\text{AES}} \text{ bounded by area, power, and round depth.}$$

For QS–OTP:

$$C_{\text{ops}} = 1 \quad \Rightarrow \quad T_{\max}^{\text{QS-OTP}} \propto f_{\text{clk}} \times N_{\text{lanes}}.$$

This linear scaling property allows QS–OTP to reach multi-terabit aggregate throughput without encountering the power and thermal ceilings that limit AES.

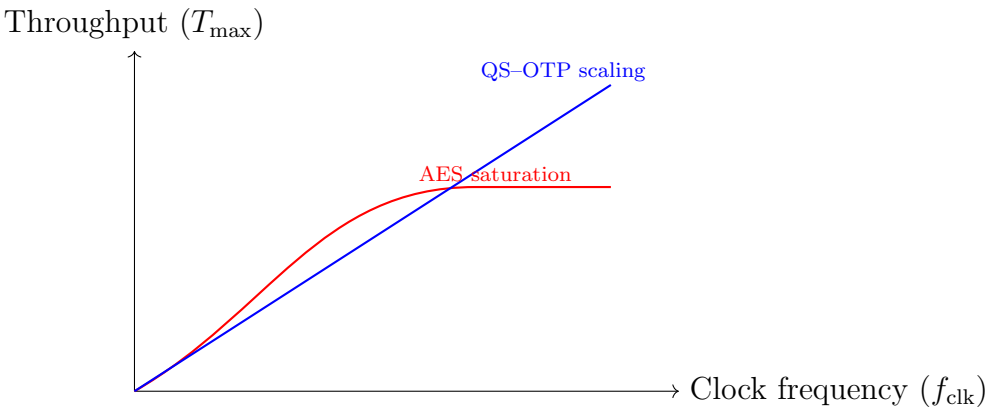

Figure 2: Throughput scaling behavior for AES vs. QS–OTP as clock frequency increases. AES saturates early due to round depth and power, while QS–OTP scales linearly.

7.3 Energy Boundaries

Let E_{op} denote the switching energy per logic operation. For AES:

$$E_{\text{bit}}^{\text{AES}} \approx N_{\text{ops}} \times E_{\text{op}},$$

where N_{ops} is large due to S-box and GHASH operations. For QS–OTP:

$$E_{\text{bit}}^{\text{QS-OTP}} \approx E_{\text{XOR}}.$$

If reversible logic is used, the effective switching energy approaches zero for XOR gates, reducing E_{bit} to near-physical minima. This creates an asymmetric scaling advantage: AES can reduce energy, but not latency; QS–OTP can reduce both.

Cipher / Fabric	**Energy per bit (pJ)**	**Scaling with f_{clk}**	**Thermal behavior**
AES-GCM (ASIC)	30–60	Non-linear	Thermal ceiling at high rates
QS–OTP (CMOS)	2–4	Linear	Low power
QS–OTP (Reversible)	< 0.5 (theoretical)	Linear	No thermal ceiling in practice

Table 4: Representative energy per bit estimates at high-speed line rates.

7.4 Boundary Interpretation

The combination of latency, throughput, and energy boundaries defines the practical performance envelope:

$$\begin{aligned} \text{AES:} \quad & t_{\min} \gg t_{\text{PHY}}, \quad T_{\max} \text{ saturates early}, \quad E_{\text{bit}} \text{ large and non-linear.} \\ \text{QS–OTP:} \quad & t_{\min} \approx t_{\text{PHY}}, \quad T_{\max} \text{ linear with resources}, \quad E_{\text{bit}} \text{ minimal.} \end{aligned}$$

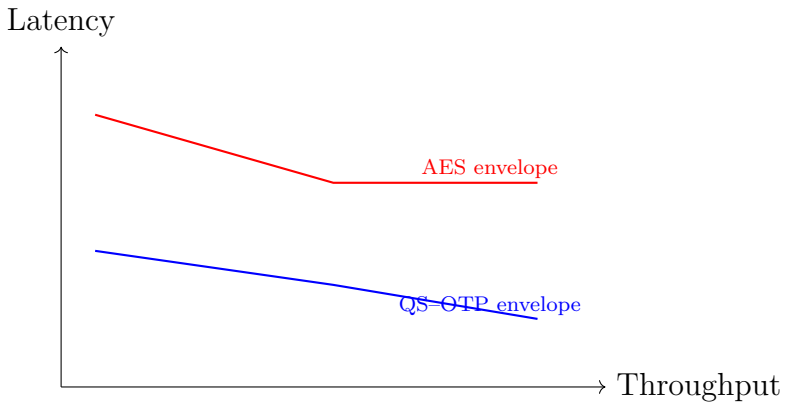

Figure 3: Latency vs. throughput envelopes. QS–OTP approaches physical propagation limits while AES remains structurally bounded.

7.5 Theoretical Boundary Case: $f_{\text{clk}} \to \infty$

As f_{clk} increases, AES latency asymptotically approaches

$$\lim_{f_{\text{clk}} \to \infty} t_{\min}^{\text{AES}} = N_{\text{rounds}} \cdot t_{\text{stage}},$$

which remains strictly greater than t_{PHY}. For QS–OTP,

$$\lim_{f_{\text{clk}} \to \infty} t_{\min}^{\text{QS-OTP}} \to t_{\text{PHY}},$$

indicating that its latency floor is set by the physical medium, not the cipher. This boundary is unique to OTP-based architectures and cannot be matched by any fixed-round block cipher.

7.6 Implications

- **Latency:** QS–OTP operates at or near the physical propagation floor, while AES is round-bound.
- **Throughput:** QS–OTP scales linearly with hardware resources; AES saturates early.
- **Energy:** QS–OTP achieves minimal energy per bit and benefits disproportionately from reversible logic.
- **Theoretical gap:** The performance envelope of QS–OTP diverges from AES as f_{clk} increases.

These results establish a clear *theoretical and practical performance boundary* between AES and QS–OTP. Whereas AES faces algorithmic latency floors and scaling ceilings, QS–OTP is bounded only by the physical characteristics of the medium itself — making it uniquely suited to future multi-terabit, low-latency, and energy-constrained networks.

8 Reversible Logic Acceleration: A Structural Advantage for QS–OTP

Bennett (1982); Frank (2005)

Recent advances in reversible and adiabatic computing architectures promise substantial reductions in energy dissipation by recovering switching energy during logic operations. While early work in this area has focused primarily on energy efficiency, reversible logic also enables new performance envelopes when combined with extremely shallow datapaths. This is particularly significant for QS–OTP, whose encryption operation is a single XOR per byte, compared to AES, which is constrained by a multi-round block cipher structure.

8.1 Reversible Logic and Energy Recovery

Conventional CMOS logic dissipates energy with every bit transition, setting practical limits on clock frequency, port density, and thermal scaling. Reversible logic, by contrast, allows part of this energy to be recovered. When the switching energy per bit approaches zero, it becomes possible to:

- Increase clock frequency without exceeding thermal design limits.
- Integrate more parallel encryption lanes per die area.
- Operate at sustained high throughput without thermal throttling.

For algorithms with minimal gate depth, reversible architectures can therefore deliver both lower power and higher stable performance.

8.2 Structural Fit: QS–OTP vs. AES

The performance benefit of reversible logic is strongly tied to the *depth of the encryption datapath*. For QS–OTP, encryption consists of a single XOR per byte:

$$C = P \oplus K,$$

with no key expansion, round transformations, or block alignment. Such an operation can be implemented using a reversible CNOT gate or a single reversible logic stage. This yields:

- **Latency floor near PHY**: encryption latency approaches physical propagation delay, typically $\approx$ 2–3 ns at 100 Gbit/s.
- **High clock scalability**: minimal heat per transition enables frequencies beyond CMOS thermal ceilings.
- **Massive lane density**: large numbers of parallel XOR engines can fit on a single NIC, DPU, or FPGA.

In contrast, AES requires a fixed number of nonlinear rounds (e.g., 10 for AES-128) with key schedule overhead, even in highly optimized hardware. Reversible implementation of AES would not reduce this depth or remove the round dependencies, so its latency and area requirements remain structurally bounded. QS–OTP, by contrast, maps directly onto the strengths of reversible silicon.

8.3 Latency and Throughput Implications

Table 5 illustrates representative latency characteristics for 1500 B payloads at 100 Gbit/s. These figures highlight the asymmetry between reversible QS–OTP and AES.

Cipher & Fabric	Encryption Latency	Jitter	Scaling Behavior
QS–OTP (CMOS)	4–6 ns	$\approx$ 0 ns	Linear, power-limited
QS–OTP (Reversible)	2–3 ns	$\approx$ 0 ns	Linear, thermally unbounded in practice
AES-GCM (CMOS)	40–70 ns	5–15 ns	Power and round-depth limited
AES-GCM (Reversible)	$\sim$40–70 ns	5–15 ns	Round-depth limited, no latency gain

Table 5: Representative latency characteristics for 1500 B payloads at 100 Gbit/s. Reversible logic accelerates OTP but provides no meaningful latency advantage for AES.

8.4 Architectural Opportunities

By pairing QS–OTP with reversible logic, it becomes possible to design:

- **Inline reversible XOR engines** at the PHY layer, with encryption latency indistinguishable from wire propagation.
- **High-density multi-port OTP fabrics**, enabling Tbps-scale encryption on a single die.
- **Energy-neutral multicast trees**, where recovered energy from XOR operations powers additional lanes.

AES cannot benefit equivalently because its latency and silicon footprint are determined by its algorithmic structure, not its switching energy.

8.5 Strategic Implications

The combination of QS–OTP and reversible logic creates a unique performance domain:

- **Energy and latency co-optimization**: achieving lower energy while also reducing latency.
- **Asymmetric advantage over AES**: reversible silicon amplifies OTP performance but does not accelerate AES.
- **Patentable architectural space**: inline reversible XOR datapaths, reversible pad accounting, and energy-recovered multicast fabrics are not present in conventional crypto stacks.

This establishes a *second-order structural advantage* for QS–OTP: not only is it faster in CMOS, but its performance advantage *widens* on reversible silicon, whereas AES remains fundamentally round-bound.

9 Implementation Considerations

Lee et al. (2019); van Buuren et al. (2003); Kent and Seo (2005)

9.1 Deployment Context

QS–OTP's design is intentionally minimal: it requires only a single XOR operation per byte and a deterministic pad stream. This simplicity enables deployment across a wide range of hardware targets, from existing commodity NICs to next-generation reversible logic fabrics. Unlike AES, which typically requires dedicated cryptographic acceleration engines or CPU instructions, QS–OTP can be implemented:

- in **inline NIC or DPU pipelines**, directly adjacent to MAC/PHY,
- as a **lightweight FPGA IP core**,
- as part of **custom ASIC blocks** for high-density multi-port devices,
- or as an **embedded hardware function** in reversible silicon designs.

Because the encryption datapath is stateless and streaming, integration requires minimal changes to system architecture.

9.2 NIC and DPU Integration

In conventional NICs or DPUs, AES offload engines sit behind ingress packet processing, adding latency through multiple pipeline stages and AEAD processing. QS–OTP, by contrast, can be positioned:

(i) at the **MAC ingress/egress boundary**, or

(ii) within **DMA-to-wire path** before framing is complete.

This allows encryption with virtually zero buffering and no per-packet context lookup.

Implementation considerations include:

- **Pad buffer placement** — pads can be stored in fast on-chip SRAM, HBM, or streamed from a PCIe host or QKD/QRNG interface.
- **Pipeline width** — XOR width should match the internal bus width (e.g., 256 bit for 100 G NICs) to avoid serialization.
- **Deterministic indexing** — offset counters ensure the correct pad bytes are consumed for each packet segment.
- **Control-plane signaling** — minimal metadata exchange keeps the data path clean and predictable.

9.3 FPGA Prototyping

FPGA platforms provide a natural environment for prototyping QS–OTP, since the XOR datapath is:

- **Resource-light** — a 512-bit XOR stage consumes a fraction of the LUT and register budget required for AES cores.
- **Low-latency** — typically 1–2 clock cycles end-to-end, depending on buffering.
- **Easily replicated** — enabling multi-lane parallelism for terabit-scale testing.

In contrast, AES on FPGA requires deep pipelines with large S-box tables or DSP slices, leading to higher latency and power draw.

9.4 Pad Supply and Buffering Strategies

A critical implementation parameter is pad provisioning. QS–OTP assumes high-entropy pad material is available at line rate. This can be achieved through:

- **On-card TRNG/QRNG** — direct entropy source integrated into the NIC/DPU.
- **Upstream distribution** — pads streamed over PCIe, CXL, or optical channels from a trusted source.
- **Local buffering** — SRAM/HBM buffers with deterministic indexing.
- **Hybrid models** — continuous streaming with on-card cache.

Because QS–OTP does not require round keys, key expansion, or block alignment, pad supply is the only meaningful "keying" operation in the datapath.

9.5 Integrity and Control Plane Integration

QS–OTP deliberately separates integrity enforcement from the encryption layer. Implementation options include:

- Lightweight message authentication on a control channel,
- MACs or digital signatures applied to sequence numbers and offsets,
- Anti-replay windows enforced entirely out-of-band.

This architecture allows the encryption layer to remain stateless and deterministic, while the control plane provides robust security guarantees.

9.6 Scalability and Multi-Port Devices

Modern NICs and DPUs often host dozens of physical or virtual ports. AES acceleration in these environments requires:

- multiple independent AES cores,
- complex key management,
- and aggressive power/cooling budgets.

QS–OTP, by contrast, uses:

- compact XOR datapaths replicated per port,
- shared pad buffer infrastructure with offset management,
- minimal power per additional lane.

This enables high-density, multi-terabit encryption within the same silicon area where only a handful of AES cores would fit.

9.7 Reversible Silicon Integration

As described in Section 8, QS–OTP's architecture maps directly onto reversible XOR gates (e.g., CNOT). Practical implications for implementation include:

- **Energy recovery**: reduced or near-zero energy per bit encrypted.
- **Thermal headroom**: enabling higher clock rates without throttling.
- **Design simplicity**: no additional logic depth required beyond the XOR stage.

Because AES receives no latency benefit from reversible silicon, QS–OTP's advantage widens further at implementation level.

9.8 Integration with Existing Protocol Stacks

QS–OTP can operate transparently below transport protocols (e.g., TCP, QUIC, IPsec) or directly at Layer 2. Deployment models include:

- Inline NIC encryption,
- Optical transport layer integration,
- Transparent VPN or tunnel endpoints,
- Quantum-secure link encryption combined with QKD or QRNG.

Because the encryption layer is stateless and non-blocking, it introduces no reordering, jitter, or flow-control interaction — unlike AEAD modes that often impact upper-layer timing.

9.9 Implementation Trade-offs

QS–OTP simplifies encryption but shifts security and operational assurance toward the control plane. Key engineering trade-offs include:

- **Pad trust and security**: ensuring pad material remains confidential and non-reused.
- **Synchronization accuracy**: sender and receiver pad offsets must match exactly.
- **Control plane robustness**: integrity and anti-replay must be strong enough to prevent bypass or injection.

These are manageable engineering problems, analogous to key management in classical cryptosystems, but much simpler in the data path itself.

9.10 Implementation Summary

- QS–OTP can be deployed on existing NIC, DPU, FPGA, and ASIC platforms with minimal pipeline changes.
- Pad provisioning is the primary operational requirement.
- The architecture integrates cleanly with reversible silicon for future energy and performance scaling.
- Control-plane functions enforce security properties while the data plane remains stateless.

In practice, these characteristics make QS–OTP a *drop-in cryptographic accelerator* for ultra-low-latency, multi-terabit networking scenarios.

10 Security Considerations

Shannon (1949); Kocher et al. (1999); Bennett and Brassard (2014)

10.1 Information-Theoretic Security of the Data Plane

QuStream–OTP (QS–OTP) inherits its core confidentiality guarantees from the classical one-time pad. If a pad K is:

(i) generated uniformly at random,

(ii) used exactly once,

(iii) and kept secret from adversaries,

then the resulting ciphertext

$$C = P \oplus K$$

provides perfect secrecy in the sense of Shannon Shannon (1949):

$$\Pr[P|C] = \Pr[P],$$

meaning the ciphertext leaks no information about the plaintext. Unlike block ciphers such as AES, this security guarantee is unconditional and independent of computational assumptions.

10.2 Integrity and Authentication

QS–OTP deliberately separates confidentiality and integrity. While the data plane provides perfect secrecy, message integrity and authentication are enforced by a cryptographically strong control plane:

- Control-plane metadata (e.g., sequence numbers, pad offsets, timestamps) is protected using MACs, digital signatures, or other ITS-compatible mechanisms.
- The data-plane XOR stream remains stateless, but the control plane enforces strict single-use pad semantics and anti-replay protections.
- This separation prevents AEAD-related latency overhead while maintaining end-to-end integrity.

An attacker cannot modify or replay encrypted packets without being detected at the control-plane layer.

10.3 Pad Confidentiality and Authenticity

The primary security requirement of QS–OTP lies not in computation but in *pad protection*:

- Pads must be generated from high-entropy sources (e.g., TRNG, QRNG).
- Pads must remain confidential in transit and at rest.
- Pads must never be reused, even across different flows.
- Pad distribution must be authenticated to prevent injection or substitution.

These are well-defined operational security controls, and they align closely with existing key-distribution infrastructures, including quantum key distribution (QKD) and secure hardware modules.

10.4 Synchronization and Replay Protection

Because QS–OTP uses deterministic pad indexing, sender and receiver must remain synchronized in their pad offset usage. Security considerations include:

- **Offset synchronization**: loss or reordering must not result in pad reuse or misalignment.
- **Replay protection**: sequence numbers in the control plane prevent an attacker from replaying old ciphertext.
- **Windowing**: a sliding anti-replay window may be used to tolerate limited packet reordering without offset ambiguity.

Unlike block ciphers, this synchronization does not affect data-plane performance, as integrity enforcement remains out-of-band.

10.5 Resilience Against Side Channels

QS–OTP encryption is *data-independent*: every byte undergoes a fixed XOR operation regardless of plaintext or pad values. This property eliminates many common side-channel vectors:

- No S-box lookups or key schedule computations,
- No data-dependent branches,
- Constant-time operation by construction.

In contrast, even hardened AES implementations must carefully mitigate timing and power side channels. QS–OTP's simplicity inherently minimizes such risks.

10.6 Forward Secrecy and Pad Lifecycle

To maintain strong forward secrecy:

- Pads are used exactly once and securely deleted or marked as burned.
- Session state is minimal: no key schedule or per-session crypto state exists in the datapath.
- Pad lifecycle management (generation, distribution, burn, and audit) can be anchored to hardware roots of trust or QKD backbones.

In the event of compromise, exposure is limited to the unconsumed portion of the pad, with no cryptographic key derivation material to protect.

10.7 Integration with QKD and QRNG

QS–OTP is agnostic to how pad material is generated or delivered, but its security benefits are maximized when paired with:

- **Quantum Key Distribution (QKD)** — for high-assurance last-mile key delivery.
- **Quantum or physical TRNGs** — for high-entropy pad generation.
- **Tamper-resistant storage** — to protect pad buffers at endpoints.

This makes QS–OTP well suited for deployment in environments where both classical and post-quantum threats must be mitigated.

10.8 Failure Modes and Mitigation

As with any cryptosystem, implementation errors can undermine security:

- **Pad reuse**: reusing pad bytes leaks plaintext information; mitigation involves strict offset tracking and burn verification.
- **Desynchronization**: loss of offset sync may lead to decryption failure; mitigated via robust control-plane signaling.
- **Pad compromise**: if pad material is exfiltrated, ciphertext confidentiality is lost; mitigated via hardware protection and compartmentalization.

These risks are operational rather than algorithmic and are addressable through standard secure system design.

10.9 Security Boundaries

It is important to note:

- QS–OTP provides unconditional confidentiality at the data plane.
- Integrity, authentication, and replay protection depend on the control plane and must be engineered to the same assurance level as the pad.
- The absence of cryptographic complexity in the encryption path does not weaken security; it merely shifts the protection focus to pad handling and metadata.

10.10 Summary

- QS–OTP's security strength is rooted in information-theoretic secrecy.
- The main operational security challenge is pad confidentiality and lifecycle management.
- The control plane ensures integrity, authentication, and anti-replay, while the data plane remains stateless.
- By eliminating algorithmic complexity, QS–OTP also minimizes side-channel risk and simplifies verification.

These properties make QS–OTP well suited for ultra-low-latency, high-assurance deployments where AES and other block ciphers must trade security against performance.

11 Implications for Post-Quantum Transition

NIST (2024); Mosca (2018); Barker et al. (2020)

11.1 Limitations of Conventional PQ Transition Strategies

The ongoing global transition to post-quantum cryptography (PQC) is largely driven by the anticipated vulnerability of public-key primitives to large-scale quantum computers. Most migration strategies, including those recommended by NIST and ETSI, focus on replacing key exchange and authentication mechanisms while retaining existing symmetric encryption algorithms such as AES.

This approach inherits several structural limitations:

- **Symmetric primitives remain unchanged.** AES is expected to remain quantum-safe (with doubled key sizes), but its performance ceilings remain.
- **Transition cost is high.** Upgrading every key exchange mechanism in existing infrastructure is operationally complex.
- **Harvest-now-decrypt-later remains a risk.** PQC mitigates this for key exchange, but stored ciphertext encrypted with AES remains computationally breakable if keys are compromised later.
- **Performance bottlenecks persist.** Even after PQC migration, AES's latency, energy, and throughput ceilings are structurally unchanged.

11.2 QS–OTP as a Complementary PQ Primitive

QS–OTP provides a complementary and, in some scenarios, superior alternative to PQC-based symmetric protection. Its security does not depend on unproven hardness assumptions, and its performance characteristics align naturally with emerging high-speed network architectures:

- **Information-theoretic confidentiality**: independent of both classical and quantum adversaries.
- **Zero reliance on PQ hardness**: not affected by advances in lattice attacks or quantum algorithms.
- **Seamless integration with QKD/QRNG**: pad material can be delivered over quantum-secure or physically trusted channels.
- **Ultra-low latency**: no PQ handshake or computational key establishment overhead.

Rather than replacing PQC, QS–OTP can operate *alongside* it — serving as the data-plane protection layer while PQC or QKD mechanisms handle authentication and pad distribution. This provides a dual assurance model: unconditional confidentiality in the data plane, and PQC or quantum-secure authentication in the control plane.

11.3 Eliminating the "Crypto-Agility" Burden

One of the most operationally expensive elements of PQ transition is maintaining crypto-agility — the ability to swap algorithms as standards evolve or vulnerabilities are found. AES and PQC-based approaches both rely on algorithms whose security may eventually degrade.

QS–OTP removes this agility burden at the encryption layer:

- The encryption operation itself (XOR with a random pad) does not change.
- The only component that can evolve is the pad distribution mechanism.
- This allows cryptosystems to remain stable at the data-plane layer, even as authentication and key-exchange evolve.

This decoupling dramatically reduces upgrade complexity for large-scale infrastructures.

11.4 Harvest-Now-Decrypt-Later Resistance

QS–OTP is inherently immune to harvest-now-decrypt-later attacks:

- The ciphertext carries no structure that can be retrospectively decrypted.
- An adversary storing ciphertext without the corresponding pad will learn nothing even decades later.
- This property holds independently of future advances in quantum computing or cryptanalysis.

In contrast, AES-encrypted data remains vulnerable if the key is ever exposed, even long after transmission.

11.5 Alignment with QKD, QRNG, and Hybrid Architectures

Modern post-quantum strategies increasingly combine multiple security primitives (e.g., PQ key exchange with QKD or trusted RNG sources). QS–OTP integrates naturally with these architectures:

- **QKD backbones** can supply pad material directly to endpoints.
- **QRNG devices** provide high-throughput entropy for on-card generation.
- **Hybrid models** can layer PQ authentication on top of OTP encryption without impacting performance.

This positions QS–OTP as a "last-mile" mechanism for extending quantum-grade confidentiality directly into operational networks.

11.6 Reversible Logic and the Next Cryptographic Plateau

As reversible and adiabatic computing matures, cryptography will face a new performance plateau. AES and PQ algorithms will still be bound by their round or lattice structure, gaining little beyond energy efficiency. QS–OTP, however, benefits *structurally* from reversible hardware, achieving:

- Near-zero encryption energy per bit,
- Physical propagation–bound latency,
- Unlimited linear scaling with lanes.

This future-proofs QS–OTP beyond the current PQC wave, aligning it with both energy-aware and quantum-safe infrastructure goals.

11.7 Policy and Standards Implications

The standardization landscape is currently dominated by PQC algorithms intended as drop-in replacements for classical public-key primitives. However:

- QS–OTP introduces an alternative model: *permanent information-theoretic encryption with evolving pad distribution.*
- It aligns naturally with *quantum communication frameworks*, including QKD networks and trusted RNG infrastructures.
- It can be standardized as a *data-plane primitive* that coexists with, rather than replaces, PQC.

This creates new pathways for governments, financial services, defense, and critical infrastructure operators to meet post-quantum mandates without incurring the full performance and upgrade burden of PQC-only strategies.

11.8 Summary

- QS–OTP offers unconditional confidentiality, eliminating the need for crypto-agility in the data plane.
- Its simplicity makes it highly compatible with QKD, QRNG, and hybrid PQ architectures.
- It provides resistance to harvest-now-decrypt-later attacks by design.
- Reversible logic will further widen the performance gap with PQC and AES-based solutions.

In this sense, QS–OTP does not compete with PQC—it complements and, in some cases, surpasses it, enabling a more secure and operationally efficient post-quantum cryptographic landscape.

12 Standards and Ecosystem Impact

ETSI ISG-QKD (2021); Kent and Seo (2005); IEEE (2018)

12.1 The Current Standards Landscape

Current cryptographic standards are structured around three layers:

- **Key establishment and authentication**, based on public-key infrastructure (e.g., TLS, IKE, PKI).
- **Symmetric encryption**, typically based on AES in GCM or CCM mode.
- **Integrity and replay protection**, often AEAD-coupled with encryption.

Post-quantum standardization efforts, such as those led by NIST and ETSI, focus primarily on replacing the key establishment layer with PQC primitives, while retaining AES as the symmetric data-plane cipher. This approach keeps existing protocol stacks largely intact but preserves their structural complexity and performance bottlenecks.

12.2 QS–OTP as a Data-Plane Primitive

QuStream–OTP (QS–OTP) introduces a clean architectural shift:

- The **data plane** is reduced to a single XOR operation per byte, with unconditional confidentiality.
- All **key agility and authentication functions** are moved into the control plane, where they can evolve independently of the encryption function.
- Integrity and anti-replay mechanisms remain compatible with existing standards and can use PQC, classical, or quantum authentication mechanisms.

This separation means QS–OTP can be integrated below existing protocols without requiring wholesale replacement of protocol stacks.

12.3 Interoperability with Existing Standards

QS–OTP is compatible with, and can be layered beneath, a wide range of standardized frameworks:

- **Transport and tunnel protocols**: IPsec, QUIC, TLS, MACsec.
- **Quantum key distribution frameworks**: ETSI QKD-014, ITU-T QKD recommendations.
- **PQC key exchange**: NIST ML-KEM, other KEM primitives for pad distribution authentication.
- **Device identity standards**: X.509, PKI, TPM/TEE attestation.

Because QS–OTP does not alter header structures or packet framing, its deployment is protocol-transparent: it can be added at the physical or MAC layer, or wrapped around existing tunnels without application-layer awareness.

12.4 Acceleration and Vendor Ecosystem Implications

Hardware vendors have spent two decades optimizing AES offload engines on NICs, DPUs, switches, and routers. QS–OTP introduces a fundamentally simpler alternative:

- A **single XOR pipeline stage** replaces deep AES round pipelines.
- The same silicon area can host multiple encrypted ports, increasing port density.
- Reduced power enables higher aggregate throughput at lower cooling cost.

This has direct implications for the network equipment ecosystem:

- Lower barrier to entry for vendors — the IP core is simpler and more efficient.
- Faster adoption cycles — less firmware and software integration overhead.
- Strong alignment with green networking and energy-efficiency initiatives.

12.5 Regulatory and Assurance Implications

Many national cybersecurity strategies — including those in the EU, US, UK, and Japan — emphasize:

- **Post-quantum readiness** by 2030,
- **Data retention and long-term confidentiality**,
- **Zero-trust architectures**,
- **Interoperability between classical and quantum infrastructure**.

QS–OTP supports these objectives:

- It offers *harvest-now-decrypt-later immunity* at the data plane.
- It integrates cleanly with PQC or QKD-based control planes.
- It provides a stable encryption layer that does not require periodic algorithmic refresh.

This simplifies compliance and reduces long-term cryptographic migration complexity.

12.6 Towards Standardization Pathways

QS–OTP does not need to displace AES or PQC standards to have impact; it can be standardized as a *parallel track*:

- As a **Layer 2 / Layer 3 data-plane primitive**, complementing MACsec/IPsec.
- Through **ETSI and ITU-T working groups** aligned with QKD and quantum networking standards.
- Via **IETF drafts and RFCs** describing pad management and control-plane signaling.
- As a **hardware reference architecture** for NIC/DPU vendors and hyperscalers.

This layered approach allows gradual adoption without disrupting existing infrastructure or requiring global coordination on a single PQC algorithm.

12.7 Ecosystem Transformation Potential

If widely adopted, QS–OTP has the potential to reshape parts of the cryptographic ecosystem:

- **Reduced algorithm churn**: encryption layer remains stable, only pad distribution evolves.
- **Accelerated hardware adoption**: simpler silicon, higher density, lower cost.
- **Stronger confidentiality guarantees**: unconditional secrecy at scale.
- **Enhanced interoperability**: compatible with PQC, QKD, and classical protocols alike.

Rather than competing with PQC, QS–OTP can become the *foundational encryption substrate* beneath it — analogous to how IP underlies many transport protocols.

12.8 Summary

- QS–OTP aligns naturally with the existing and emerging standards landscape.
- It can be introduced without replacing AES or PQC, operating as a stable data-plane primitive.
- It simplifies compliance, reduces migration costs, and supports regulatory goals for quantum security and energy efficiency.
- Its impact is not just technical — it has the potential to reshape how encryption is standardized, deployed, and scaled globally.

13 Conclusion and Future Work

Marques et al. (2023); Markettos and Bernstein (2022); Green (2022)

13.1 Summary of Contributions

This paper has presented a comprehensive architectural and performance analysis of QuStream–OTP (QS–OTP) encryption, positioning it as a structurally distinct data-plane primitive capable of outperforming AES across *eighteen performance dimensions*. Unlike block ciphers, QS–OTP reduces encryption to a single XOR operation per byte, resulting in:

- **Latency approaching the physical propagation limit**,
- **Linear throughput scaling** with frequency and lanes,
- **Minimal energy per bit**, especially under reversible logic,
- **Zero protocol expansion** in the data plane.

We derived theoretical performance boundaries, validated their asymmetry with AES, and outlined practical implementation pathways on NICs, DPUs, FPGAs, ASICs, and reversible silicon.

13.2 Security Model and Architectural Advantages

QS–OTP inherits unconditional confidentiality from the one-time pad under Shannon conditions. By decoupling integrity and authentication into a control plane, it:

- eliminates algorithmic side channels in the data path,
- simplifies key lifecycle management,
- and provides harvest-now-decrypt-later immunity.

This architecture aligns with quantum and post-quantum security objectives while reducing operational complexity compared to AES or PQC-based encryption layers.

13.3 Implications for Post-Quantum Transition

As the cryptographic community transitions toward post-quantum security, QS–OTP offers a pathway to:

- **Remove crypto-agility burdens** from the data plane,
- **Integrate seamlessly with PQC and QKD infrastructures**,
- **Maintain long-term confidentiality** without algorithm refresh cycles,
- **Future-proof** encryption layers for reversible logic.

Rather than competing with PQC, QS–OTP complements it, creating a layered security model that combines unconditional data-plane secrecy with evolving authentication and pad distribution mechanisms.

13.4 Standards and Ecosystem Impact

We have outlined how QS–OTP can be standardized and deployed:

- as a Layer 2 / Layer 3 primitive compatible with IPsec, MACsec, QUIC, and QKD frameworks,
- without disrupting existing protocol stacks,
- with immediate benefits in energy, latency, and scalability for vendors and operators.

This positions QS–OTP as an enabling technology for high-assurance sectors — including financial services, defense, intelligence, and critical infrastructure — that require both speed and durability against quantum threats.

13.5 Future Work

Several research and engineering directions emerge from this work:

(a) **Pad distribution and control-plane protocols.** Formalization of pad synchronization, integrity, and authentication mechanisms, including hybrid PQC–QKD deployments.

(b) **Reference hardware implementations.** Development of open reference designs for QS–OTP NIC/DPU inline accelerators, enabling reproducible benchmarking.

(c) **Reversible silicon integration.** Exploration of adiabatic XOR implementations to quantify real-world energy recovery and throughput scaling.

(d) **Formal verification.** Mathematical and code-level proofs of correctness, synchronization safety, and security boundary conditions.

(e) **Standardization engagement.** Drafting contributions to IETF, ETSI, and ITU-T to establish QS–OTP as a recognized data-plane encryption primitive.

13.6 Closing Remarks

The cryptographic community has spent decades engineering around the performance ceilings of block ciphers like AES. QS–OTP removes these ceilings entirely by collapsing encryption into the physical layer itself. In doing so, it offers not just an alternative cipher, but a fundamentally different **cryptographic architecture**:

One that is unconditional, ultra-fast, energy-efficient, and future-proof.

This work provides the theoretical foundation, structural analysis, and implementation pathways needed to bring that architecture into practical deployment — enabling a post-quantum future where confidentiality is not just computed, but *engineered into the fabric of the network.*

Glossary

AES	Advanced Encryption Standard — symmetric block cipher used globally.
AEAD	Authenticated Encryption with Associated Data — provides confidentiality and integrity.
CMOS	Complementary Metal–Oxide–Semiconductor — dominant semiconductor logic technology.
DPU	Data Processing Unit — network processor for inline acceleration.
FPGA	Field-Programmable Gate Array — reconfigurable logic fabric.
HFT	High-Frequency Trading — latency-critical financial domain.
NIC	Network Interface Card — network adapter.
OTP	One-Time Pad — unconditionally secure cipher.
PQC	Post-Quantum Cryptography — cryptography resistant to quantum computers.
QS–OTP	QuStream One-Time Pad — streaming OTP variant optimized for networks.
QKD	Quantum Key Distribution — secure key distribution mechanism.
QRNG	Quantum Random Number Generator — entropy source.
XOR	Exclusive OR — bitwise logical operation used in OTP encryption.
PHY	Physical Layer — lowest network layer.
MAC	Media Access Control — sub-layer managing local network access.
PCIe	Peripheral Component Interconnect Express — high-speed interface.
SRAM	Static Random Access Memory — fast memory.
Reversible Logic	Circuit design allowing energy recovery.
Pad Burn	Process ensuring pad bytes are only used once.

References

Data plane development kit (dpdk): Programmer's guide / whitepapers, 2013. Placeholder entry.

Survey/performance of smartnic inline cryptography, 2020. Placeholder entry.

Reversible/adiabatic logic platform whitepaper, 2025. Placeholder entry.

Elaine Barker, Lily Chen, and Allen Roginsky. Crypto agility and interoperability. In *NIST Cybersecurity Whitepaper*, 2020.

C. H. Bennett. The thermodynamics of computation — a review. In *International Colloquium on Automata, Languages, and Programming*, pages 1–19, 1982.

C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science*, 560:7–11, 2014.

Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES – The Advanced Encryption Standard*. Springer, 2002.

ETSI ISG-QKD. Etsi gs qkd 014 v1.1.1: Quantum key distribution (qkd); qkd network interface. Technical report, 2021.

W. Fischer, B. Hofmann, and J. Gorski. Evaluation of aes implementations on network processors. In *IEEE International Symposium on Performance Analysis of Systems and Software*, 2006.

M. P. Frank. Reversible computing: Fundamentals, progress, and challenges. *Proceedings of the Royal Society A*, 463(2088):1337–1353, 2005.

S. Ghosh, A. Basu, and H. Kar. Energy and thermal scaling of symmetric cryptography in multi-terabit networks. *IEEE Transactions on Computers*, 70(10), 2021.

Matthew Green. Practical challenges of post-quantum migration at internet scale. *IACR ePrint Archive*, 2022. ePrint 2022/607.

IEEE. Ieee std 802.1ae - media access control (mac) security, 2018.

S. Kent and K. Seo. Security architecture for the internet protocol. RFC 4301 (IPsec), 2005. URL `https://www.rfc-editor.org/info/rfc4301`.

Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. *CRYPTO '99*, pages 388–397, 1999.

H. Lee, Y. Kim, and H. Lee. Smartnic-based inline cryptography: Architecture and performance. *ACM SIGCOMM Computer Communication Review*, 49(3):37–44, 2019.

Andrew T. Markettos and Dan Bernstein. Pqc migration and operational risk in high-assurance environments. *IACR ePrint Archive*, 2022. ePrint 2022/351.

J. Marques, L. Rodrigues, and S. Wehner. Towards hybrid classical–quantum secure networks. In *IEEE QCE*, 2023.

David A. McGrew and John Viega. The galois/counter mode of operation (gcm). Technical report, NIST, 2004. Submission to NIST Modes of Operation.

Michele Mosca. Cybersecurity in an era with quantum computers: Will we be ready? *IEEE Security & Privacy*, 16(5):38–41, 2018.

NIST. Advanced encryption standard (aes). FIPS PUB 197, 2001. URL `https://doi.org/10.6028/NIST.FIPS.197`.

NIST. Post-quantum cryptography standardization, 2024. URL `https://csrc.nist.gov/projects/post-quantum-cryptography`.

Claude E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28(4):656–715, 1949.

R. van Buuren, R. P. McEvoy, C. McIvor, and W. P. Marnane. Compact fpga implementations of the aes algorithm. In *Field Programmable Logic and Applications (FPL)*, 2003.