Privacy-Enhancing Technologies (PET) Comparative Ecosystems and Supply Chain Risk Analysis for FluidCluster Models:

Architectural Topology, Cryptographic Primitives (FHEVM, TEE, MPC, PSI, ZKP), and Institutional Hybridization Across Layers 0–3

FluidCluster Sovereign Liquidity & Cryptography Research Group
research@fluidcluster.pt | https://www.fluidcluster.pt

November 2026 — Working Technical Blueprint (v1.0)

Abstract

[SKELETON PLACEHOLDER — 250 words]
Context: Enterprise ERP ecosystems, open banking rails (PSD3/PSR), and multilateral B2B clearing networks face an acute structural paradox: solving systemic liquidity gridlocks and supply-chain Bullwhip distortions requires global graph visibility, yet exposing plaintext commercial ledgers violates trade secrets and regulatory mandates.
Objective: This paper establishes a multi-layered supply-chain topology (Layers 0 to 3) comparing Fully Homomorphic Encryption (FHE/fhEVM) and Threshold MPC-KMS architectures against Trusted Execution Environments (TEEs) and hybrid cryptographic stacks (MPC, PSI, ZKP) instantiated in the *FluidCluster / Ecompensa* architecture.
Methodology: Combining Design Science Research Methodology (DSRM) with Socio-Technical Supply Chain Risk Management (SCRM) and STRIDE-PET threat modeling, we evaluate silicon fabrication bottlenecks, patent/licensing barriers, and regulatory friction across 40+ ecosystem actors.
Contributions: (1) A five-tier architectural taxonomy (Layers 0, 1, 2, 2.5, 3); (2) comparative risk matrices contrasting hardware root-of-trust (TEE) with mathematical post-quantum blindness (FHE); and (3) a formal migration roadmap toward Collaborative FHE-ERPs and Photonic-Mutual Banks.

Keywords: Fully Homomorphic Encryption (FHE), fhEVM, Trusted Execution Environments (TEE), Multi-Party Computation (MPC), Private Set Intersection (PSI), Zero-Knowledge Proofs (ZKP), Supply Chain Risk Management (SCRM), FluidCluster, PSD3, Silicon Photonics.

Contents

Architectural Diagram & Figure Specification Registry **4**

1 Introduction, Theoretical Framework, and Hypotheses **5**
1.1 The Enterprise Coordination & Liquidity Paradox 5
1.2 Consensual Research Methodology: DSRM, STS-SCRM, and STRIDE-PET . . . 5
1.3 Core Hypotheses and Logical Validation Protocol 6

2 Cryptographic Paradigms: FHE/fhEVM and Threshold MPC-KMS vs. TEEs **7**
2.1 Trusted Execution Environments (TEEs) & Hardware-Rooted Trust 7
2.2 Fully Homomorphic Encryption (FHE), fhEVM, and Threshold MPC-KMS . . . 8
2.3 Auxiliary Primitives in FluidCluster Hybrids (MPC, PSI, ZKP) 9

2.4 Comparative Evaluation Matrix: TEE vs. FHEVM vs. Hybrid PETs 10
2.5 Formal Composability & Security Boundary Analysis 10

3 Ecosystem Architecture and Player Topology (Layers 0 through 3) 12
3.1 Layer 0: Hardware Acceleration (GPUs, FPGAs, ASICs, Opto-Photonic PICs) . 14
3.2 Layer 1: Cryptographic Primitives, Compilers, & Enterprise Enclave Engines . . 15
3.3 Layer 2: Confidential Smart Contracts, fheDAOs, & Off-Chain Coprocessors . . . 16
3.4 Layer 2.5: Middleware, API Connectors, Agentic MCP Swarms, & Auditing Oracles 17
3.5 Layer 3: Traditional Finance, Open Banking (PSD3/PSR), & Enterprise ERPs . 17

4 Supply Chain Topology Risk Analysis and Inter-Layer Barriers 18
4.1 Layer 0 $\leftrightarrow$ Layer 1: Silicon Geopolitics, Photonic Packaging, & ISA Fragmentation 19
4.2 Layer 1 $\leftrightarrow$ Layer 2: Technical, Cryptographic, & IP/Patent Barriers 20
4.3 Layer 2 $\leftrightarrow$ Layer 2.5: Coprocessor Centralization, Oracle Collusion, & Swarm Risks 20
4.4 Layer 2.5 $\leftrightarrow$ Layer 3: Legal, Regulatory, & Compliance Friction 21

5 Ecosystem Verticalizations and Layer 3 Institutional Hybridization 22
5.1 Verticalization Vectors Across the PET Stack 22
5.2 Hybrid Archetype I: The Collaborative FHE-Based ERP 23
5.3 Hybrid Archetype II: The Photonic-Mutual Bank (FHEVM Sovereign Mutual) . 23
5.4 Migration Roadmap for Incumbent Banks, PSPs, and ERP Vendors 25

6 Synthesis, Risk Scoring Matrix, and Future Research Agenda 26
6.1 Hypothesis Validation Summary (H_1, H_2, H_3) 26
6.2 Composite Supply Chain Risk Scorecard Across Architectures 26
6.3 Open Problems and Next-Generation Research Directions 28

List of Figures

1 Cryptographic Security vs. Computational Latency Frontier (H_1 and H_2 validation): illustrating the Trust-Performance Inversion between pure TEEs and monolithic FHE, and how FluidCluster's hybrid pipeline reaches the institutional Pareto target zone. 9
2 Layer 0 Hardware Acceleration Supply Chain Risk Graph: contrasting foundry dependencies, thermal dissipation, and fabrication bottlenecks across GPUs, FPGAs, custom FHE ASICs, and Opto-Photonic PICs. 12
3 Master Five-Tier Privacy-Enhancing Technology (PET) Supply Chain Topology ($L_0 \rightarrow L_3$), mapping hardware accelerators, cryptographic compilers, confidential smart contracts, AI agentic middleware, and regulated institutional endpoints. . 13
4 Inter-Layer Friction & Barrier Matrix ($L_0 \leftrightarrow L_3$): cross-referencing technical bottlenecks, intellectual property/foundry concentration, and European statutory constraints alongside FluidCluster architectural mitigations. 19
5 FluidCluster Blind Triangle & Off-Chain Coprocessor Execution Loop: decoupling off-chain PSI topology discovery and client-side ZK-SNARK validation from `fhEVM` consensus, photonic-accelerated Simplex netting, and Threshold MPC-KMS / PSD3 settlement. 24
6 Layer 3 Institutional Hybridization: architectural coupling between the *Collaborative FHE-Based ERP* (left) and the *Photonic-Mutual Bank* (right), uniting local ERP privacy, photonic-accelerated `fhEVM` clearing, Threshold MPC-KMS regulatory auditability, and PSD3 PISP/AISP settlement. 29

List of Tables

1 Comparative Architectural Matrix of Privacy-Enhancing Technologies (PETs) Across Trust, Performance, and Institutional Criteria 11
2 Composite Supply Chain & STRIDE-PET Risk Scorecard Across Enterprise Clearing Architectures 27

Architectural Diagram Figure Specification Registry

- **Figure 1 (Master Topology):** *The Five-Tier PET Supply Chain Stack (Layers 0, 1, 2, 2.5, 3)* — Mapping silicon/photonic hardware, cryptographic compilers, fhEVM chains, MCP AI swarms, and PSD3/ERP institutional endpoints.
- **Figure 2 (Threat Trust Inversion):** *Cryptographic Security vs. Computational Latency Frontier* — Plotting pure TEE, Threshold MPC, ZKP-Rollups, Pure TFHE/CKKS, and FluidCluster Hybrid (Off-Chain PSI + ZKP + fhEVM + Photonic Coprocessor).
- **Figure 3 (Layer 0 Silicon Photonic Bottlenecks):** *Hardware Acceleration Supply Chain Risk Graph* — Foundry concentration (TSMC/Samsung), export controls, and architectural divergence between GPUs (NVIDIA H100/B200), FPGAs (AMD/Xilinx), ASICs (Niobium/SEMIFIVE, Cornami, Chain Reaction, Fabric), and Opto-Photonic PICs (Optalysys MZI Fourier cores).
- **Figure 4 (FluidCluster Blind Triangle Coprocessor Loop):** *End-to-End Execution Pipeline in FluidCluster 2.0/3.0* — Illustrating local ERP ingestion, client-side ZK-SNARK commitments, off-chain PSI loop discovery, Leontief/Simplex multilateral netting on fhEVM (‘euint64‘/‘ERC-7984‘), and Threshold MPC-KMS settlement release.
- **Figure 5 (Inter-Layer Friction Barrier Matrix):** *Technical, IP/Patent, and Regulatory Chokepoints (L0 $\leftrightarrow$ L3)* — Visualizing BSD-3-Clause-Clear vs. proprietary patents, FHE noise/bootstrapping latency, and PSD3/DORA/GDPR compliance boundaries.
- **Figure 6 (Layer 3 Hybridization Evolution):** *From Siloed Legacy ERPs to Collaborative FHE-ERPs and Photonic-Mutual Banks* — Architectural transition from centralized plaintext clearinghouses to sovereign cryptographic mutuals.

1 Introduction, Theoretical Framework, and Hypotheses

1.1 The Enterprise Coordination & Liquidity Paradox

Modern enterprise networks, open banking infrastructures, and multilateral business-to-business (B2B) clearing systems operate within an acute structural contradiction: *systemic liquidity optimization requires global graph visibility, yet commercial survival and regulatory compliance mandate strict local data confidentiality* [3, 10]. In conventional corporate trade networks, mutual obligations form dense, cyclic directed graphs of accounts payable and accounts receivable. When bilateral invoices remain siloed inside isolated Enterprise Resource Planning (ERP) databases, firms cannot detect closed debt cycles ($A \to B \to C \to A$). Consequently, enterprises suffer artificial working-capital gridlocks, forcing reliance on costly short-term credit lines or factoring intermediaries even when the network possesses sufficient net offsetting claims to settle internally [2].

Simultaneously, opacity across multi-tier supply chains amplifies demand-signal distortion—widely documented as the *Bullwhip Effect*—where minor fluctuations in downstream retail orders trigger escalating inventory oscillations upstream [10]. While sharing real-time order books, inventory buffers, and invoice graphs across tiers would mathematically dissolve both liquidity traps and Bullwhip variance, exposing plaintext commercial ledgers to a centralized clearinghouse or peer counterparties introduces severe adversarial risks:

1. **Margin Counterparty Extraction:** Revealing supplier identities, unit pricing, and payment terms enables dominant buyers or clearing operators to disintermediate suppliers or capture monopoly rents.

2. **Systemic Cyber-Attack Surface:** Centralizing plaintext financial telemetry creates high-value honeypots. Empirical security audits indicate that approximately 73% of mid-market ERP installations lag critical security patches by more than six months, leaving monolithic API connectors exposed to credential harvesting and lateral traversal [9].

3. **Statutory Regulatory Breach:** Uncontrolled cross-entity disclosure of transaction graphs violates commercial trade-secret protections, banking secrecy statutes, and European data protection mandates (GDPR), while failing operational resilience thresholds under the EU Digital Operational Resilience Act (DORA) [10, 34].

Privacy-Enhancing Technologies (PETs)—specifically Trusted Execution Environments (TEEs), Fully Homomorphic Encryption (FHE and fhEVM), Threshold Multi-Party Computation Key Management Systems (MPC-KMS), Private Set Intersection (PSI), and Zero-Knowledge Proofs (ZKPs)—promise to resolve this paradox by enabling computation over encrypted or attested state [19, 20]. However, transitioning PETs from isolated cryptographic primitives to institutional financial infrastructure reveals that security, latency, and compliance are governed by an intricate, five-tier industrial supply chain spanning silicon/photonic hardware (Layer 0) to regulated payment service providers and ERPs (Layer 3).

1.2 Consensual Research Methodology: DSRM, STS-SCRM, and STRIDE-PET

To evaluate confidential computing architectures rigorously without falling into single-layer cryptographic reductionism, this study adopts a consensual, tri-axial research methodology combining:

1. **Design Science Research Methodology (DSRM):** Following Hevner et al. [33], we treat the *FluidCluster / Ecompensa* architecture [1, 6] as an instantiated socio-technical artifact designed to solve multilateral B2B netting and supply-chain coordination. DSRM

structures our iterative build-and-evaluate loop, testing how specific architectural refinements—such as offloading cyclic graph discovery to off-chain PSI [5] or accelerating polynomial multiplications via Mach-Zehnder Interferometer (MZI) photonic cores [4]—affect system-wide feasibility.

2. **Socio-Technical Systems Supply Chain Risk Management (STS-SCRM):** Drawing on multi-tier resilience frameworks [31, 32], we model the PET ecosystem as a vertically coupled supply chain ($L_0 \to L_1 \to L_2 \to L_{2.5} \to L_3$). Rather than assessing cryptographic protocols in isolation, STS-SCRM tracks how upstream disruptions—such as semiconductor foundry concentration, proprietary patent walls, or cloud enclave deprecations—propagate downstream to impair institutional solvency and regulatory compliance.

3. **STRIDE-PET Adversarial Threat Modeling:** We extend classical STRIDE threat taxonomy (*Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege*) to cryptographic pipelines. Specifically, we distinguish between *cryptographic confidentiality* (protection of ciphertext payloads under Ring Learning With Errors, RLWE) and *topological metadata privacy* (resistance against transaction-graph inference, timing correlation, and access-control list leakage on public or permissioned ledgers [15]).

Across all comparative evaluations in subsequent chapters, architectures are benchmarked against five formal state variables:

- **Trusted Computing Base Surface ($\mathcal{S}_{\mathbf{TCB}}$):** The minimal set of hardware, firmware, and software components whose compromise invalidates confidentiality or integrity.
- **Ciphertext Expansion Bandwidth Overhead ($\Omega_{\mathbf{BW}}$):** The ratio of encrypted wire/storage footprint to underlying plaintext representation ($\text{Size}(c)/\text{Size}(m)$).
- **End-to-End Clearing Latency ($\tau_{\mathbf{clear}}$):** Wall-clock execution time from ERP invoice ingestion to cryptographic settlement finality, decomposed into client proof generation (τ_{zkp}), topology discovery (τ_{psi}), homomorphic evaluation (τ_{fhe}), and threshold decryption (τ_{mpc}).
- **Vendor Foundry Lock-In Entropy ($H_{\mathbf{lock}}$):** Degree of dependency on single-source silicon foundries, proprietary instruction sets, or closed-source patents.
- **Statutory Auditability Index ($\mathcal{A}_{\mathbf{reg}}$):** Capacity to prove compliance with PSD3/PSR, DORA, and tax-authority invariants (e.g., SAF-T/AT invoice authenticity) without revealing sensitive commercial parameters to unauthorized observers [8, 29].

1.3 Core Hypotheses and Logical Validation Protocol

Grounded in preliminary empirical telemetry from the FluidCluster testbed and current 2026 industry benchmarks, we formulate three falsifiable hypotheses to guide our multi-layer analysis:

Hypothesis H_1 (The Trust-Performance Inversion):
Statement: Pure Trusted Execution Environment (TEE) architectures minimize computational latency ($\tau_{\text{clear}} < 1$ ms) by anchoring security in hardware-enforced memory encryption, but suffer from an irreducible Trusted Computing Base ($\mathcal{S}_{\text{TCB}}$) exposed to microarchitectural side-channels and silicon foundry supply-chain choke points. Conversely, pure Fully Homomorphic Encryption (FHE) eliminates hardware trust assumptions via post-quantum lattice hardness, but incurs 10^3–$10^5\times$ ciphertext expansion and bootstrapping overhead that render unaccelerated on-chain clearing economically unviable without specialized Layer 0 ASIC or opto-photonic acceleration [2, 14, 22].
Logical Validation Protocol: Tested in Chapters 2, 3, and 4 by contrasting enclave attestation boundaries (Intel TDX, AMD SEV-SNP, NVIDIA Confidential Compute) against

TFHE/CKKS bootstrapping complexity across GPUs, FPGAs, custom ASICs (Niobium, Cornami), and photonic Fourier cores (Optalysys).

Hypothesis H_2 (Modular Hybrid Convergence):
Statement: Orchestrating stateful encrypted smart contracts (`fhEVM` at Layer 2) with off-chain Private Set Intersection (PSI) for cycle discovery, client-side ZK-SNARKs for input well-formedness, and TEE-isolated Agentic Swarms operating via Model Context Protocol (MCP at Layer 2.5) resolves the privacy-latency-verifiability trilemma. Specifically, decoupling graph pruning (PSI) from homomorphic optimization (fhEVM) reduces encrypted gate evaluations by orders of magnitude while preventing on-chain metadata leakage [5, 6, 19].
Logical Validation Protocol: Tested in Chapters 2 and 4 using FluidCluster's Blind Triangle execution pipeline, evaluating throughput scaling ($> 50,000$ TPS target envelope) and STRIDE-PET resilience against metadata correlation and threshold key-holder collusion.

Hypothesis H_3 (Layer 3 Institutional Hybridization):
Statement: Emerging European regulatory regimes—specifically the Third Payment Services Directive and Payment Services Regulation (PSD3/PSR), DORA, and statutory multilateral set-off frameworks (e.g., Portuguese Decree-Law no. 150/2019)—create structural friction for public, unpermissioned blockchains while simultaneously penalizing centralized plaintext data aggregation. Consequently, incumbent Layer 3 actors (ERPs, commercial banks, PISPs, and AISPs) will evolve into hybrid socio-technical architectures—specifically *Collaborative FHE-ERPs* and *Photonic-Mutual Banks*—that separate fiduciary liquidity settlement from plaintext data custody [9, 8, 29]. *Logical Validation Protocol:* Tested in Chapter 5 by mapping statutory PSD3/PSR open-banking interfaces (ASPSP/PISP/AISP) and DORA ICT third-party risk mandates onto two formal institutional reference architectures.

Chapter 1 Explainer: Thesis Core Methodological Compass

Core Thesis: No single Privacy-Enhancing Technology (PET) satisfies enterprise throughput, post-quantum cryptographic blindness, and statutory financial auditability simultaneously; achieving institutional viability requires a vertically integrated, hybrid cryptographic supply chain. **Main Takeaways:**

- **The Coordination Paradox:** Businesses need network-wide visibility to net mutual debts and prevent supply-chain Bullwhip distortions, yet exposing raw invoices destroys competitive margins and violates banking secrecy.
- **Beyond Protocol Reductionism:** Evaluating FHE or TEEs solely on mathematical or CPU benchmarks ignores fatal supply-chain dependencies, such as semiconductor foundry concentration at Layer 0 and PSD3 compliance rules at Layer 3.
- **Three Testable Pillars:** Our framework tests whether hardware-accelerated hybrid pipelines (PSI + ZKP + fhEVM + MCP Swarms) can outperform monolithic TEE or FHE designs and catalyze new hybrid institutions like Collaborative FHE-ERPs and Photonic-Mutual Banks.

2 Cryptographic Paradigms: FHE/fhEVM and Threshold MPC-KMS vs. TEEs

2.1 Trusted Execution Environments (TEEs) & Hardware-Rooted Trust

Trusted Execution Environments (TEEs) enforce confidentiality and code integrity through hardware-level isolation inside modern CPUs and GPUs, including Intel Trust Domain Exten-

sions (TDX), AMD Secure Encrypted Virtualization-Secure Nested Paging (SEV-SNP), ARM Confidential Compute Architecture (CCA), and NVIDIA Hopper/Blackwell Confidential Computing [27, **?**]. By instantiating hardware-encrypted memory regions (enclaves or confidential virtual machines) governed by dedicated on-die memory encryption engines, TEEs execute standard compiled binaries over plaintext data at near-native clock speeds, incurring minimal computational overhead ($\tau_{\text{clear}} < 1$ ms) [20].

To establish trust with external counterparties, TEEs rely on cryptographic *remote attestation*. Upon initialization, the processor generates a hardware-signed measurement (cryptographic hash) of the enclave's initial memory state and firmware revision, anchored in vendor-burned fuses (e.g., Intel Provisioning Certification Keys or AMD Versioned Chip Endorsement Keys). Key Management Services (KMS) verify this attestation quote before releasing decryption keys into the protected memory space.

Despite superior throughput and developer ergonomics, TEEs exhibit an irreducible Trusted Computing Base ($\mathcal{S}_{\text{TCB}}$) that introduces three structural vulnerabilities for sovereign financial clearing:

1. **Microarchitectural Side-Channels & Fault Injection:** Because enclaves share physical silicon caches, branch predictors, and power rails with untrusted host hypervisors, adversaries with physical or co-resident access can extract secret keys via cache-timing attacks, speculative execution flaws (Spectre/Meltdown descendants, SGAxe), or dynamic voltage frequency scaling (Plundervolt) [19].

2. **Foundry & Vendor Root-of-Trust Centralization:** Security rests entirely on the integrity of the silicon manufacturer's attestation root keys and microcode updates, creating single-vendor lock-in (H_{lock}) and geopolitical supply-chain exposure.

3. **Absence of Public Verifiability:** While remote attestation proves *what binary ran*, it does not produce a succinct, non-interactive mathematical proof that external auditors or smart contracts can verify independently without trusting the hardware vendor [19, 20].

2.2 Fully Homomorphic Encryption (FHE), fhEVM, and Threshold MPC-KMS

Fully Homomorphic Encryption (FHE) eliminates hardware trust assumptions altogether by enabling arbitrary arithmetic and boolean circuits to execute directly over ciphertexts. Grounded in the hardness of the Ring Learning With Errors (RLWE) problem, FHE provides post-quantum cryptographic blindness: an untrusted cloud server, blockchain validator, or photonic coprocessor evaluates functions without ever possessing the decryption key [1, 11].

Modern enterprise and Web3 deployments bifurcate across three complementary RLWE scheme families:

- **TFHE (Torus FHE):** Optimized for fast *programmable bootstrapping* (≤ 1 ms per gate on modern GPUs [14]), enabling exact boolean gates, encrypted integer comparisons (`FHE.lt`, `FHE.eq`), and oblivious multiplexers (`FHE.select`) essential for conditional smart-contract logic [12].

- **BFV / BGV:** Leveled schemes supporting Single Instruction, Multiple Data (SIMD) packing for high-throughput exact integer arithmetic.

- **CKKS:** Approximate fixed- and floating-point arithmetic over complex/real vectors, uniquely suited for homomorphic linear programming, Leontief input-output matrices, and Network Simplex optimization in multilateral debt clearing [4, 2].

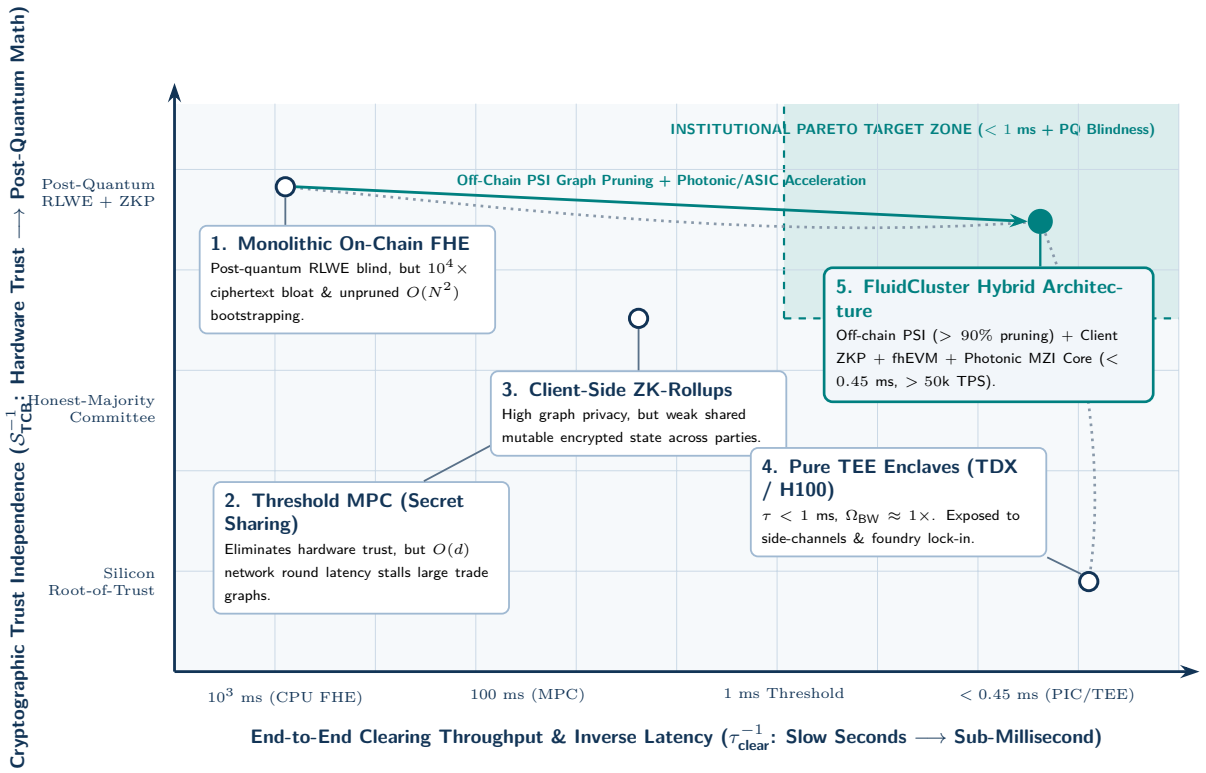


Figure 1: Cryptographic Security vs. Computational Latency Frontier (H_1 and H_2 validation): illustrating the Trust-Performance Inversion between pure TEEs and monolithic FHE, and how FluidCluster's hybrid pipeline reaches the institutional Pareto target zone.

Symbolic Execution in `fhEVM`: Running heavy polynomial Number Theoretic Transforms (NTT) directly inside Ethereum Virtual Machine (EVM) consensus would stall block production. Consequently, the `fhEVM` architecture decouples consensus from homomorphic computation via *symbolic execution* [12, 13]. Smart contracts manipulate lightweight 256-bit ciphertext handles (`euint64`, `ebool`) and emit deterministic cryptographic execution events. A network of off-chain FHE coprocessors listens to these events, executes the heavy TFHE/CKKS bootstrapping asynchronously, and commits updated ciphertext pointers alongside verifiable execution attestations. Fine-grained data visibility is enforced on-chain via programmable Access Control Lists (`FHE.allow`, `FHE.allowTransient`), permitting selective decryption by authorized counterparties or statutory auditors [12, 15].

Threshold MPC-KMS: A global FHE network requires a single shared public key so ciphertexts from disparate enterprises can interact homomorphically. To prevent any single infrastructure operator from holding the corresponding secret decryption key, `fhEVM` integrates a *Threshold Multi-Party Computation Key Management System (MPC-KMS)* [11]. Using Distributed Key Generation (DKG), the master secret key sk is partitioned into n shares across an independent committee of KMS nodes such that any coalition of fewer than t nodes learns zero information about sk. Decryption requests—triggered only when on-chain ACL rules are satisfied—require a t-out-of-n threshold signature and partial decryption aggregation, neutralizing single-point-of-compromise threats.

2.3 Auxiliary Primitives in FluidCluster Hybrids (MPC, PSI, ZKP)

While FHE provides universal encrypted compute, evaluating global $N \times N$ corporate invoice graphs blindly inside FHE incurs prohibitive $O(N^2)$ ciphertext complexity and risks leaking

topological metadata (e.g., observing which encrypted accounts interact on-chain) [5, 15]. To establish empirical validation for Hypothesis H_2, the *FluidCluster / Ecompensa* architecture composes three specialized auxiliary PETs around the `fhEVM` core:

1. **Off-Chain Private Set Intersection (PSI):** Prior to on-chain submission, participating ERPs execute multi-party PSI protocols over hashed invoice counterparty identifiers [5]. PSI isolates only those subgraphs forming closed liquidity cycles ($A \to B \to C \to A$) while keeping non-intersecting commercial relationships completely hidden from both the clearing engine and non-participating firms. This pre-filtering prunes sparse enterprise graphs by $> 90\%$, drastically reducing downstream FHE circuit depth.

2. **Client-Side Zero-Knowledge Proofs (ZK-SNARKs / zkVMs):** Because FHE guarantees computation integrity *after* encryption but cannot verify whether a user's encrypted input represents a valid, tax-compliant obligation, FluidCluster requires local ERP connectors to generate client-side ZK-SNARK proofs [6, 16]. These proofs attest that: (a) the encrypted receivable matches an official tax-authority signed electronic invoice (e.g., Portuguese AT/SAF-T hash), (b) values are non-negative within valid bit-widths, and (c) the ciphertext is well-formed under the network's FHE public key.

3. **Collaborative MPC for Game-Theoretic Settlement:** In multi-party corporate restructuring, M&A coalitions, and joint procurement modules, lightweight secret-sharing MPC protocols compute cooperative Shapley value allocations and anti-gaming invariants before committing final net deltas to `ERC-7984` confidential ciphertokens [7, 36].

2.4 Comparative Evaluation Matrix: TEE vs. FHEVM vs. Hybrid PETs

Table 1 synthesizes the structural trade-offs across pure TEEs, Threshold MPC, ZK-Rollups, Monolithic FHE, and the FluidCluster Hybrid stack across our six evaluation dimensions.

2.5 Formal Composability & Security Boundary Analysis

A critical insight of our Design Science evaluation is that hybrid PET pipelines do not simply sum the guarantees of their constituent primitives; rather, security depends on the *cryptographic handoff interfaces* between layers [19, 16]. In the FluidCluster architecture, cross-primitive composability is governed by three formal isolation boundaries:

1. **Client-to-Ledger Verifiable Encryption Bridge (ZKP $\to$ fhEVM):** Standard RLWE encryption does not prevent a malicious participant from submitting malformed ciphertexts or negative invoice amounts to distort multilateral netting. By binding a client-side ZK-SNARK proof π_{in} to the ciphertext commitment $c = \text{Enc}_{pk}(x)$ via shared Pedersen/Poseidon hash commitments [6], the smart contract rejects invalid inputs at the consensus boundary without learning x.

2. **Topological Decoupling Boundary (PSI $\to$ fhEVM):** Executing Private Set Intersection off-chain ensures that only counterparties sharing an active, closed debt cycle submit ciphertexts to the shared `fhEVM` state [5]. Furthermore, batching cycle settlements into zero-sum net vectors masks individual invoice weights against on-chain graph-inference attacks [2, 15].

3. **Enclave-to-Ciphertext Containment (TEE Swarm $\leftrightarrow$ Threshold MPC-KMS):** Where autonomous AI agents (Layer 2.5 MCP Swarms) perform semantic invoice normalization or anomaly scoring inside hardware TEEs [35], they operate strictly under time-bounded, revocable Access Control List grants (`FHE.allowTransient`) [12]. Even if a physical TEE enclave suffers a side-channel compromise, the adversary gains access only to ephemeral, single-batch micro-context rather than the master Threshold MPC-KMS decryption key sk.

Table 1: Comparative Architectural Matrix of Privacy-Enhancing Technologies (PETs) Across Trust, Performance, and Institutional Criteria

Evaluation Dimension	**Pure TEE Enclaves** (TDX / SEV-SNP / H100)	**Threshold MPC** (Secret Sharing)	**ZK-Rollups** (Client-Side SNARKs)	**Pure On-Chain FHE** (Monolithic TFHE)	**FluidCluster Hybrid** (PSI + ZKP + fhEVM + PIC)
Root of Trust ($\mathcal{S}_{\mathbf{TCB}}$)	Hardware silicon fuses & vendor firmware attestation	Honest-majority / threshold non-collusion ($< t$ nodes)	Mathematical soundness (elliptic curves / hashes)	Mathematical RLWE lattice hardness + Threshold KMS	Dual-anchored: RLWE math + threshold KMS + client ZKPs
Post-Quantum Resistance	Vulnerable (classical ECDSA/RSA attestation roots)	Protocol-dependent (lattice/hash variants exist)	Mixed (Groth16/Plonk vulnerable; STARKs PQ-safe)	**Natively PQ-Safe** (RLWE lattice assumptions)	**PQ-Safe Core** (TFHE/CKKS + lattice-compatible commitments)
Shared Mutable Encrypted State	Native (high-speed plaintext inside enclave RAM)	High communication round complexity ($O(d)$ rounds)	Weak (prover must know all plaintext inputs)	**Native** (multi-user computation on shared ciphertexts)	**Native + Pruned** (PSI isolates active cycles first)
Latency & Throughput ($\tau_{\mathbf{clear}}$)	$<$ 1 ms (near-native CPU/GPU throughput)	10–500 ms (network-bound bandwidth)	Fast verify ($<$ 10 ms); slow client proving (1–10 s)	High latency (10^3–$10^5\times$ overhead on CPU)	$<$ 0.45 ms core solver (Photonic/GPU coprocessor)
Metadata & Graph Privacy	Vulnerable to host memory/network traffic analysis	Protected if pairwise channels encrypted	High (shielded UTXO pools hide sender/receiver)	Moderate (ACL handles leak interaction topology)	**High** (Off-chain PSI + batch zero-sum netting)
Regulatory Auditability ($\mathcal{A}_{\mathbf{reg}}$)	Code attestation only; opaque runtime state	Requires online committee quorum per audit	Selective disclosure via viewing keys / ZK proofs	Granular on-chain ACL (`FHE.allow`) for regulators	**Programmable** (ACL + ZK-AML + SAF-T tax proofs)

Chapter 2 Explainer: Cryptographic Trade-Offs in Plain Terms

Core Thesis: TEEs act like a *locked armored vault* (fast computation inside, but you must trust the vault manufacturer), whereas FHE acts like *manipulating contents inside a sealed glovebox* (the operator never sees the data, guaranteed by post-quantum lattice math).

Main Takeaways:

- **Why Pure TEEs Fall Short:** Near-zero computational overhead makes TEEs ideal for orchestrating AI agents, yet hardware side-channel vulnerabilities or malicious foundry insiders can compromise enclave memory.
- **How fhEVM + Threshold MPC-KMS Works:** Smart contracts manipulate lightweight ciphertext pointers while off-chain coprocessors perform the heavy math, and splitting the master key across a threshold committee ensures no single operator can unilaterally decrypt balances.
- **The FluidCluster Hybrid Advantage:** Using PSI to find debt loops off-chain, ZKPs to prove invoice validity on the client, and fhEVM to net balances on-chain avoids wasting expensive FHE operations on non-intersecting corporate data—while strict cryptographic handoff boundaries contain any potential enclave breach.

3 Ecosystem Architecture and Player Topology (Layers 0 through 3)

To evaluate systemic dependencies, vendor lock-in (H_{lock}), and cross-layer performance bottlenecks, we decompose the confidential computing and sovereign clearing landscape into a five-tier supply chain topology ($L_0 \to L_1 \to L_2 \to L_{2.5} \to L_3$). Figure 3 illustrates how physical photonic and silicon compute primitives at the base propagate up to regulated open-banking endpoints and ERP ledgers at the institutional edge.

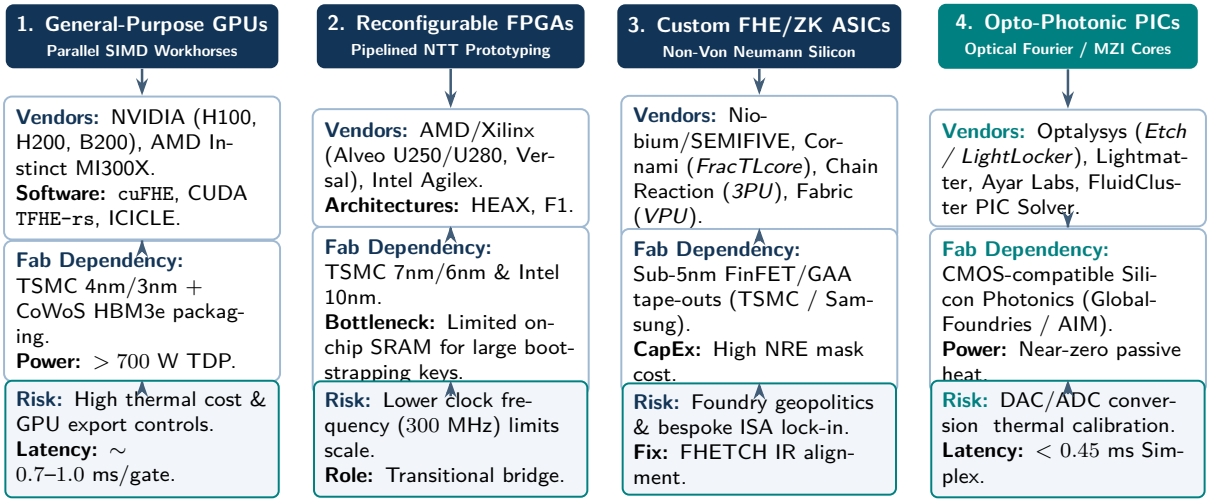


Figure 2: Layer 0 Hardware Acceleration Supply Chain Risk Graph: contrasting foundry dependencies, thermal dissipation, and fabrication bottlenecks across GPUs, FPGAs, custom FHE ASICs, and Opto-Photonic PICs.

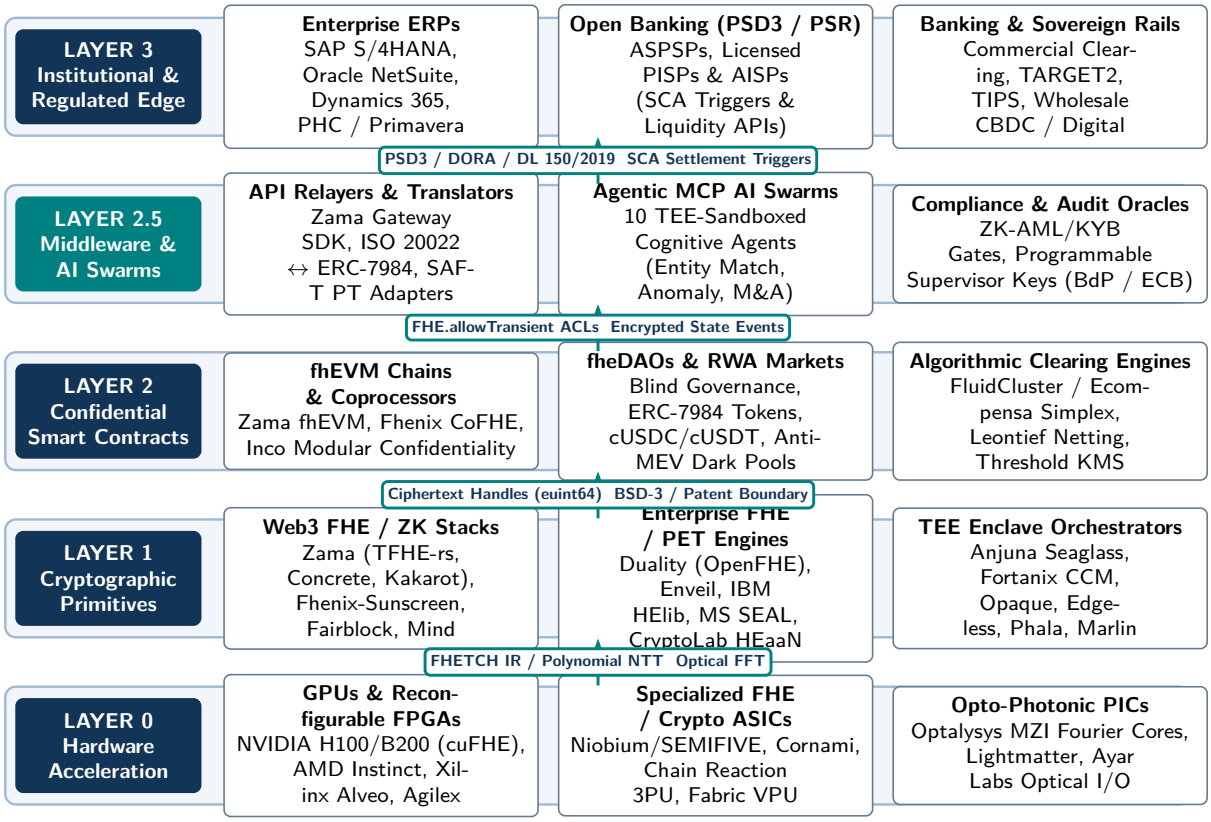


Figure 3: Master Five-Tier Privacy-Enhancing Technology (PET) Supply Chain Topology ($L_0 \rightarrow L_3$), mapping hardware accelerators, cryptographic compilers, confidential smart contracts, AI agentic middleware, and regulated institutional endpoints.

3.1 Layer 0: Hardware Acceleration (GPUs, FPGAs, ASICs, Opto-Photonic PICs)

The computational bottleneck of Fully Homomorphic Encryption lies in high-degree polynomial multiplication over cyclotomic rings ($N \in [2^{10}, 2^{16}]$) executed via Number Theoretic Transforms (NTT) and Fast Fourier Transforms (FFT), alongside frequent programmable bootstrapping [1, 14]. Layer 0 comprises four distinct hardware acceleration classes competing to reduce bootstrapping latency from milliseconds to sub-microsecond regimes:

1. **General-Purpose Parallel Silicon (GPUs):** Commercial data-center GPUs—led by **NVIDIA** (H100, H200, and Blackwell B200 architectures utilizing `cuFHE` and Zama's CUDA-accelerated `TFHE-rs` backend) and **AMD** (Instinct MI300X)—currently serve as the baseline production workhorse for `fhEVM` coprocessors, achieving single-gate TFHE bootstrapping in < 1 ms [14]. However, GPUs suffer from memory-bandwidth walls when streaming multi-megabyte bootstrapping keys and incur high thermal power dissipation (> 700 W per accelerator).

2. **Reconfigurable Gate Arrays (FPGAs):** Field-Programmable Gate Arrays—specifically **AMD/Xilinx** (Alveo U250/U280 and Versal Adaptive SoCs) and **Intel/Altera** (Agilex series powering academic/industrial architectures such as HEAX and F1)—provide custom pipelined NTT butterfly units and on-chip High Bandwidth Memory (HBM) routing without multi-million-dollar mask fabrication costs. FPGAs serve primarily as transitional prototyping platforms for cloud FHE enclaves.

3. **Specialized FHE & Cryptographic ASICs:** Purpose-built Application-Specific Integrated Circuits eliminate von Neumann instruction-fetch overhead through spatial dataflow architectures and massive on-chip SRAM banks:

   - **Niobium Microsystems:** Emerged from the DARPA DPRIVE (Data Protection in Virtual Environments) program; partnered with **SEMIFIVE** in 2026 to tape out custom high-throughput FHE accelerator ASICs targeting cloud and defense workloads [39, 23].
   - **Cornami:** Pioneer of a scalable, non-von Neumann reconfigurable computational fabric (*FracTLcore*) paired with its *SCIFR* compiler API, engineered for massively parallel, low-power full-circuit FHE evaluation [24, 25].
   - **Chain Reaction:** Developer of the *Electrum* and *3PU* (Privacy-Preserving Processing Unit) custom ASICs tailored for real-time FHE cloud acceleration, and co-founding member of the FHETCH hardware consortium [21].
   - **Fabric Cryptography:** Architect of the *Verifiable Processing Unit (VPU)*, a specialized cryptographic processor unifying ZKP multi-scalar multiplication (MSM) and FHE polynomial primitives on a single die.
   - **Ingonyama & Intel HERACLES: Ingonyama** (developer of the *ICICLE* cryptographic acceleration library bridging GPUs and custom cryptographic silicon [**?**]) alongside **Intel's HERACLES** (DARPA DPRIVE) and **Duality/SRI** hardware co-design initiatives.

4. **Opto-Photonic Integrated Circuits (PICs):** To overcome fundamental electronic Joule-heating and RC interconnect capacitance limits in electronic NTT/FFT execution, opto-photonic accelerators compute Fourier transforms passively at the speed of light:

   - **Optalysys:** Leader in silicon-photonic FHE acceleration (*Etch* and *LightLocker* platforms), utilizing integrated Mach-Zehnder Interferometer (MZI) meshes and optical Fourier transform cores to execute polynomial multiplications with near-zero thermal dissipation [22, 21].

- **Lightmatter & Ayar Labs:** Providers of photonic tensor cores (*Envise*) and co-packaged optical I/O interconnects (*TeraPHY*) that eliminate chip-to-chip memory bandwidth bottlenecks when streaming large RLWE ciphertexts.
- **FluidCluster Photonic Solver Integration:** Within the FluidCluster 2.0/3.0 architecture, MZI photonic cores execute homomorphic CKKS/TFHE Simplex pivoting and Leontief matrix inversions in < 0.45 ms, sustaining $> 50,000$ TPS clearing throughput [1, 4].

3.2 Layer 1: Cryptographic Primitives, Compilers, & Enterprise Enclave Engines

Layer 1 abstracts raw Layer 0 silicon and photonic instructions into developer-usable cryptographic libraries, circuit compilers, and hardware attestation runtime engines. The Layer 1 ecosystem bifurcates into three distinct industrial cohorts:

1. **Web3 & Open-Source FHE Protocol Stacks:**
   - **Zama:** Creator of the `TFHE-rs` Rust library, the `Concrete` ML/FHE compiler, and the `fhEVM` confidential smart-contract protocol; expanded its zero-knowledge and rollup interoperability stack via the 2026 acquisition of **Kakarot ZK** [11, 13].
   - **Fhenix:** Developer of the `CoFHE` off-chain FHE coprocessor architecture for Ethereum and Arbitrum; strengthened its compiler and ZKP-FHE hybrid capabilities through its July 2026 acquisition of **Sunscreen** (creator of the *Parasol* FHE compiler) [17].
   - **Inco Network:** Provider of a modular *Universal Confidentiality Layer* combining `fhEVM` encrypted state with TEE-accelerated confidential compute nodes to offer flexible latency-security profiles across EVM chains [18].
   - **Mind Network, Fairblock, & Octra: Mind Network** specializes in FHE restaking and zero-trust HTTPZ communication layers; **Fairblock** (*FairyRing*) integrates Identity-Based Encryption (IBE), Threshold MPC, and FHE for pre-execution MEV protection; **Octra** develops hypergraph-based FHE execution environments.
2. **Enterprise, Defense, & Big-Tech FHE/PET Compilers:**
   - **Duality Technologies:** Co-founder of the open-source **OpenFHE** library (unifying PALISADE, HElib, and FHEW); commercializes federated encrypted SQL, secure multi-party analytics, and privacy-preserving AI for regulated banking and healthcare consortia [26].
   - **Enveil:** Pioneer of *ZeroReveal* encrypted search and cross-boundary intelligence analytics built on homomorphic primitives for defense and financial crime investigations.
   - **Big-Tech & Academic Reference Libraries: IBM Research** (`HElib` and `HElayers` for encrypted neural networks), **Microsoft Research** (`SEAL` BFV/CKKS engine), **Google** (HE Transpiler and *Private Join and Compute* combining PSI with homomorphic encryption), and **CryptoLab** (`HEaaN` approximate CKKS optimization engine).
3. **Enterprise TEE & Confidential Computing Orchestrators:**
   - **Anjuna Security:** Developer of the *Seaglass* platform enabling zero-code "lift-and-shift" deployment of unmodified enterprise applications into Intel TDX, AMD SEV-SNP, and NVIDIA GPU enclaves [28].
   - **Fortanix:** Provider of the *Confidential Computing Manager (CCM)* and *Data Security Manager (DSM)*, integrating hardware-enforced TEE orchestration with enterprise HSM key governance [27].

- **Opaque Systems, Edgeless Systems, & R3 Conclave: Opaque** delivers confidential Spark/SQL analytics and encrypted LLM inference; **Edgeless Systems** (*Constellation*, *Contrast*) provides confidential Kubernetes clusters; **R3 Conclave** targets regulated multi-party financial workflows.
- **Decentralized TEE Coprocessors: Phala Network** and **Marlin** (*Oyster*) operate verifiable, blockchain-anchored TEE worker networks used extensively for Layer 2.5 AI agent sandboxing [20].

3.3 Layer 2: Confidential Smart Contracts, fheDAOs, & Off-Chain Coprocessors

Layer 2 transforms raw Layer 1 cryptographic primitives into stateful, multi-party decentralized applications, confidential financial markets, and algorithmic clearing engines. Rather than forcing every validator node to execute resource-intensive homomorphic bootstrapping, modern Layer 2 networks adopt a modular separation between *consensus state* and *off-chain cryptographic coprocessing*:

1. **Chain & Symbolic Coprocessor Topologies:**
   - **Symbolic Coprocessor Rollups:** Exemplified by Zama's `fhEVM` Coprocessor (deployed across Ethereum L1 and Layer 2 rollups such as Base) and Fhenix's `CoFHE` coprocessor on Arbitrum [13, 17]. Smart contracts store 256-bit ciphertext handles (`euint64`, `eaddress`) while specialized coprocessor networks execute asynchronous TFHE/CKKS graphs and post threshold-signed results back to the host chain.
   - **Modular Confidentiality Overlays:** Networks such as Inco Network provide confidentiality-as-a-service to external EVM chains via cross-chain messaging bridges and dual FHE/TEE execution lanes [18].
2. **fheDAOs & Confidential Web3 Financial Markets:**
   - **Blind Sovereign Governance (`fheDAOs`):** Decentralized Autonomous Organizations utilizing encrypted ballot tallying and homomorphic quadratic/token-weighted voting to eliminate whale coercion, last-minute bandwagoning, and governance front-running [37].
   - **Confidential Token Standards & Dark Pools:** Implementation of `ERC-7984` confidential fungible tokens—supporting encrypted balances and encrypted transfers—enables institutional dark pools, anti-MEV automated market makers, and tokenized Real-World Assets (RWAs) such as encrypted corporate receivables and confidential stablecoins (`cUSDC`, `cUSDT`) [36, 8].
3. **Algorithmic Off-Chain Clearing Engines (*FluidCluster / Ecompensa*):**
   - Operating as an institutional Layer 2 algorithmic engine, FluidCluster pairs `fhEVM` settlement contracts with off-chain mathematical optimization solvers (Pyomo/HiGHS Network Simplex and Leontief input-output solvers) accelerated by Layer 0 GPUs and optophotonic MZI cores [2, 4].
   - Through *Vectorized Zero-Sum Obfuscation*, cyclic trade debts ($A \to B \to C \to A$) are netted homomorphically so only minimal net residual deltas hit the ledger, preserving complete commercial blindness across all participants [3, 6].

3.4 Layer 2.5: Middleware, API Connectors, Agentic MCP Swarms, & Auditing Oracles

A major barrier to institutional PET adoption is the impedance mismatch between deterministic, ciphertext-native Layer 2 smart contracts and messy, asynchronous Layer 3 enterprise records. Layer 2.5 bridges this gap through three specialized middleware functions:

1. **API Integrators, Relayers, & Schema Translators:**
   - Cryptographic relayer SDKs (e.g., Zama Gateway/Relayer) abstract client-side key management, ciphertext re-encryption, and EIP-712 signature bundling for enterprise webhooks.
   - Domain-specific schema adapters map legacy banking messages (**ISO 20022** `pain.001` / `camt.053`) and European e-invoicing standards (EN 16931, Portuguese **SAF-T PT**) into zero-knowledge commitments and `ERC-7984` ciphertoken minting payloads [9, 8].
2. **AI-Assisted MPC/TEE Providers (Agentic Swarms):**
   - **FluidCluster Cognitive Swarm:** An orchestrated federation of 10 specialized autonomous AI agents communicating via the **Model Context Protocol (MCP)** inside hardware-attested TEE enclaves (Intel TDX / NVIDIA H100 Confidential Compute) [35, 1].
   - Operating under revocable, time-bounded `FHE.allowTransient` permissions, these agents resolve semantic entity mismatches across heterogeneous supplier SKUs, detect fraudulent circular invoicing anomalies, and negotiate Pareto-optimal debt-to-equity fusion proposals (*FluidCluster M&A Aperture*) before sealing final parameters back into FHE ciphertexts [6, 7].
3. **Cryptographic Auditing & Regulatory Compliance Oracles:**
   - Decentralized oracle networks and zero-knowledge compliance gates verify Know-Your-Business (KYB), Anti-Money Laundering (AML), and Counter-Terrorist Financing (CTF) attestations without publishing corporate ownership trees on-chain.
   - Programmable regulatory viewing keys allow supervisory bodies (e.g., Banco de Portugal, ECB, or national tax authorities) to audit systemic liquidity ratios and statutory tax withholdings on demand [34, 8].

3.5 Layer 3: Traditional Finance, Open Banking (PSD3/PSR), & Enterprise ERPs

Layer 3 encompasses the incumbent financial, regulatory, and corporate systems where commercial obligations originate and final fiat settlement occurs:

1. **Enterprise Resource Planning (ERP) Systems:**
   - Global Tier-1 suites (**SAP S/4HANA**, **Oracle NetSuite**, **Microsoft Dynamics 365**) alongside regional SME platforms (e.g., **PHC**, **Primavera** across Iberia) serve as the primary system of record for accounts payable, accounts receivable, and supply-chain inventory [9].
   - In the FluidCluster topology, lightweight local ERP sidecar connectors extract hashes of tax-validated invoices, execute off-chain PSI cycle discovery, and generate client-side ZK-SNARKs without exposing raw relational database tables externally [5, 6].
2. **Commercial Banks, Clearing Houses, & Sovereign Settlement Rails:**

- Tier-1 commercial banks, national clearing houses, and central bank settlement infrastructures (**TARGET2**, **TIPS**, and wholesale Central Bank Digital Currency / Digital Euro pilots) provide ultimate fiat finality [8].
- Multilateral FHE netting reduces gross interbank settlement volume by eliminating self-canceling circular trade debts prior to invoking central bank liquidity rails.

3. **Open Banking & Payment Service Providers (PSD2 → PSD3 / PSR):**

- Under the European Union's revised **Third Payment Services Directive (PSD3)** and directly applicable **Payment Services Regulation (PSR)** (which merge traditional Payment Institutions and E-Money Institutions into a unified licensing regime), financial intermediaries bifurcate into three operational roles [29, 30]:
  - **ASPSPs (Account Servicing Payment Service Providers):** Credit institutions holding corporate current accounts, mandated to provide high-availability, obstacle-free dedicated Open Banking APIs and user permission dashboards.
  - **PISPs (Payment Initiation Service Providers):** Regulated entities authorized to trigger automated credit transfers (e.g., SEPA Instant) directly from a corporate debtor's bank account upon receiving a threshold-signed multilateral netting settlement instruction.
  - **AISPs (Account Information Service Providers):** Regulated aggregators authorized to read real-time account balances and liquidity telemetry to verify solvency thresholds prior to admitting an enterprise into an FHE clearing cycle.

Chapter 3 Explainer: Mapping the 5-Tier Value Chain

Core Thesis: Confidential computing has evolved from isolated cryptographic libraries into a specialized five-tier industrial supply chain stretching from silicon-photonic wafers (Layer 0) to PSD3-regulated bank APIs and corporate ERPs (Layer 3).
Main Takeaways:

- **Hardware-to-Compiler Foundation (Layers 0–1):** Optical Fourier chips (Optalysys) and custom ASICs (Niobium, Cornami) supply the raw polynomial throughput required by compiler stacks (Zama, Fhenix-Sunscreen, Duality) and enclave engines (Anjuna, Fortanix).
- **Encrypted Execution & AI Middleware (Layers 2–2.5):** Confidential blockchains (`fhEVM`) rely on off-chain algorithmic coprocessors (FluidCluster) and TEE-isolated AI swarms (MCP agents) to normalize messy corporate invoices into verified encrypted optimization problems.
- **Institutional Endpoints (Layer 3):** Legacy ERPs (SAP, Oracle, Primavera) and PSD3-regulated PISPs/AISPs sit at the top edge, acting as the cryptographic data-origination and residual fiat-settlement rails.

4 Supply Chain Topology Risk Analysis and Inter-Layer Barriers

Applying our Socio-Technical Supply Chain Risk Management (STS-SCRM) and STRIDE-PET frameworks reveals that mathematical guarantees at the protocol layer can be severely degraded—or entirely invalidated—by physical hardware dependencies below them and statutory compliance regimes above them [31, 32]. Figure 4 maps the primary technical, proprietary, and regulatory chokepoints across the four inter-layer boundaries.

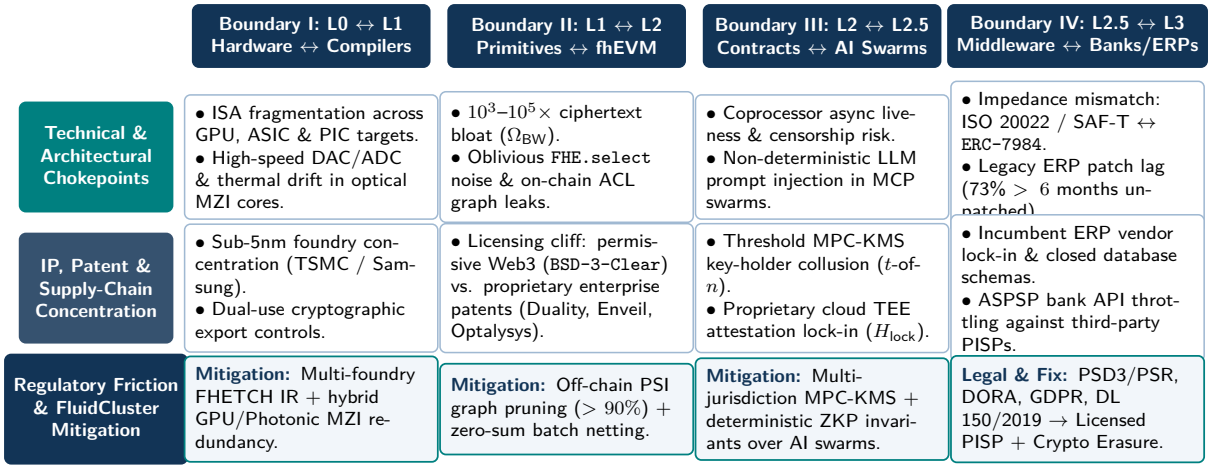


Figure 4: Inter-Layer Friction & Barrier Matrix ($L_0 \leftrightarrow L_3$): cross-referencing technical bottlenecks, intellectual property/foundry concentration, and European statutory constraints alongside FluidCluster architectural mitigations.

4.1 Layer 0 ↔ Layer 1: Silicon Geopolitics, Photonic Packaging, & ISA Fragmentation

The interface between physical hardware accelerators (Layer 0) and cryptographic compilers/libraries (Layer 1) exhibits the highest capital expenditure barrier and vendor lock-in entropy (H_{lock}) across the entire confidential computing stack:

1. **Semiconductor Foundry Concentration & Export Controls:** High-throughput FHE ASICs (such as Niobium/SEMIFIVE, Cornami, and Chain Reaction [23, 24, 21]) as well as enclave-capable GPUs (NVIDIA H100/B200) rely overwhelmingly on advanced sub-5nm FinFET and Gate-All-Around (GAA) fabrication nodes concentrated in East Asian foundries (TSMC and Samsung Foundry). Geopolitical disruptions, wafer allocation bottlenecks, or dual-use cryptographic export restrictions directly threaten Layer 0 supply continuity for sovereign European financial infrastructure [1].

2. **Opto-Photonic Packaging & Digital-to-Analog Conversion Overhead:** While Silicon Photonics Integrated Circuits (PICs)—such as Optalysys's Mach-Zehnder Interferometer (MZI) Fourier cores [22] and FluidCluster's optical Simplex coprocessor [4]—eliminate electronic Joule heating during polynomial multiplication, they introduce severe packaging and conversion friction. Moving high-precision RLWE polynomial coefficients between electronic HBM memory and optical waveguides requires ultra-high-speed Digital-to-Analog (DAC) and Analog-to-Digital (ADC) converters. Calibration drift across thermal gradients and lower wafer-level packaging yields remain critical scaling risks compared to mature CMOS electronic dies.

3. **Instruction-Set Architecture (ISA) & Compiler Fragmentation:** Unlike standard deep learning, which converged around CUDA, PyTorch, and ONNX, FHE hardware lacks a unified Intermediate Representation (IR). Each Layer 0 vendor exposes bespoke low-level primitives—ranging from NVIDIA CUDA kernels (`cuFHE`) to Cornami's spatial *SCIFR* API [24] and Optalysys's optical Fourier instruction pipeline [22]. Although the FHE Technical Consortium for Hardware (**FHETCH**, co-founded by Optalysys, Chain Reaction, and Niobium [21]) aims to standardize hardware-software interfaces, Layer 1 compilers (`Concrete`, `OpenFHE`, `Parasol`) still require costly manual kernel retargeting per accelerator target.

4.2 Layer 1 ↔ Layer 2: Technical, Cryptographic, & IP/Patent Barriers

Translating low-level homomorphic circuit libraries (Layer 1) into multi-tenant decentralized applications, `fheDAOs`, and multilateral clearing smart contracts (Layer 2) encounters four structural friction points:

1. **Ciphertext Bloat & Network Bandwidth Amplification ($\Omega_{\mathbf{BW}}$):** A single 64-bit plaintext integer (`uint64`) expands into a multi-kilobyte RLWE ciphertext (`euint64`), yielding expansion factors of $\Omega_{\text{BW}} \in [10^3, 10^5]$ [2, 12]. Transmitting thousands of encrypted corporate invoices across peer-to-peer Layer 2 gossip networks saturates consensus bandwidth unless offloaded to external data-availability layers or pruned upstream via off-chain PSI [5].

2. **Oblivious Execution Penalty (`FHE.select`) & Bootstrapping Noise Budgets:** Because smart-contract execution cannot branch conditionally on encrypted booleans without leaking execution paths, every `if/else` branch in `fhEVM` must evaluate *both* execution paths and multiplex the result homomorphically via `FHE.select` [12]. Deeply nested loops—such as iterative Network Simplex pivoting in debt clearing—rapidly exhaust the ciphertext's RLWE noise budget, triggering expensive programmable bootstrapping operations that stall throughput unless offloaded to specialized Layer 0 coprocessors [4, 14].

3. **Topological Metadata Leakage & ACL Inference:** As demonstrated in recent cryptanalytic evaluations of public `fhEVM` deployments [15], encrypting transaction *amounts* does not conceal transaction *graphs*. Adversaries observing on-chain ciphertext handle updates and Access Control List (`FHE.allow`) permission grants can infer commercial supply-chain relationships, payment frequencies, and liquidity stress through graph-topology and timing analysis. Mitigating this leak requires batch zero-sum netting and off-chain PSI isolation [2, 5].

4. **Intellectual Property (IP) Walls & Licensing Cliffs:** A sharp licensing asymmetry divides Layer 1 from Layer 2. Web3-native FHE protocols (Zama, OpenFHE, Fhenix) rely on permissive or dual open-source licenses (such as Zama's `BSD-3-Clause-Clear` patent-grant license) to foster permissionless smart-contract composability [11, 26]. Conversely, enterprise PET pioneers (Duality, Enveil, Cornami, Optalysys) hold dense patent portfolios covering encrypted SQL joins, hardware-accelerated NTT pipelines, and optical Fourier transforms. Integrating enterprise-grade Layer 1 primitives into open Layer 2 protocols risks patent-infringement exposure or proprietary royalty cliffs.

4.3 Layer 2 ↔ Layer 2.5: Coprocessor Centralization, Oracle Collusion, & Swarm Risks

While offloading heavy homomorphic computation and AI-assisted invoice normalization to Layer 2/2.5 coprocessors resolves on-chain gas bottlenecks, it introduces three socio-technical trust boundaries that must be hardened under STRIDE-PET threat modeling:

1. **Off-Chain Coprocessor Liveness & Censorship Risks:** In symbolic `fhEVM` architectures, smart contracts emit ciphertext execution events that depend on external coprocessor nodes to compute and return valid handles [13]. If the coprocessor set is overly centralized or lacks cryptographic verifiable-computation proofs (such as succinct ZK proofs of correct TFHE bootstrapping), a compromised or offline coprocessor can censor specific corporate clearing batches or halt settlement liveness. FluidCluster mitigates this risk via multi-operator coprocessor redundancy paired with deterministic cross-checking [1].

2. **Threshold MPC-KMS Key-Holder Collusion:** Because all participants encrypt invoices under a shared network FHE public key, the t-out-of-n Threshold MPC-KMS commit-

tee represents the ultimate cryptographic choke point [11]. If $\geq t$ key-share custodians collude—or suffer simultaneous supply-chain compromise of their underlying cloud HSMs—they could covertly decrypt historical ledger ciphertexts without triggering on-chain ACL events. Mitigating this requires institutional jurisdiction diversity (distributing shares across central banks, clearing houses, and audited cloud enclaves) combined with proactive key-share refreshment and economic slashing [3].

3. **Adversarial Poisoning & Hallucination in Agentic MCP Swarms:** Deploying autonomous AI agents (Layer 2.5 Model Context Protocol swarms) inside TEEs to normalize heterogeneous ERP invoices or propose M&A debt-to-equity fusions introduces non-deterministic AI failure modes [35, 7]. Malicious counterparties could embed adversarial prompt injections inside invoice metadata fields to manipulate agentic risk scoring. FluidCluster neutralizes this threat by stripping unstructured text upstream and subjecting every swarm output to deterministic, non-AI smart-contract invariants (e.g., strict conservation of zero-sum net flow and verified client ZK-SNARK signatures) before settlement execution [6].

4.4 Layer 2.5 $\leftrightarrow$ Layer 3: Legal, Regulatory, & Compliance Friction

The upper boundary connecting cryptographic middleware (Layer 2.5) to incumbent banks, Payment Service Providers, and corporate ERPs (Layer 3) is governed by strict statutory frameworks that frequently clash with permissionless blockchain assumptions:

1. **EU PSD3 & Payment Services Regulation (PSR) Friction:** Under the 2026 EU PSD3/PSR framework, Account Servicing Payment Service Providers (ASPSPs) must expose dedicated Open Banking APIs with granular user permission dashboards, while Payment Initiation Service Providers (PISPs) face strict Strong Customer Authentication (SCA) and Authorized Push Payment (APP) fraud liability rules [29, 30]. Automated smart-contract settlement triggers cannot legally initiate fiat SEPA Instant transfers unless mediated by a licensed Layer 2.5/3 PISP bridge that maps multi-party cryptographic authorizations to compliant SCA tokens.

2. **Digital Operational Resilience Act (DORA) ICT Risk Mandates:** DORA requires European banks and financial entities to maintain comprehensive register-of-information audits, exit strategies, and penetration-testing regimes over all critical Information and Communication Technology (ICT) third-party service providers [34]. Relying on anonymous, permissionless `fhEVM` coprocessors or opaque foreign cloud TEE providers violates DORA concentration and auditability rules unless Layer 2.5 enforces permissioned, SLA-backed institutional node whitelisting.

3. **GDPR Erasure vs. Ledger Immutability & Statutory Set-Off Law:** Storing encrypted corporate or sole-proprietor invoice records on an immutable ledger creates tension with the GDPR "Right to Erasure" (Article 17) if ciphertexts are ever rendered decryptable by future cryptanalytic breakthroughs. FluidCluster resolves this via *Cryptographic Erasure*: personal identifiers remain strictly off-chain inside local ERPs, while on-chain ciphertext access is governed by revocable ACLs and ephemeral session keys [15]. Simultaneously, multilateral clearing operations must satisfy national commercial set-off statutes—such as Portuguese **Decree-Law no. 150/2019** governing multilateral invoice compensation—requiring deterministic, tax-auditable cryptographic proofs of mutual consent and invoice authenticity (SAF-T PT) [10, 38].

Chapter 4 Explainer: Where the Supply Chain Breaks

Core Thesis: Cryptographic perfection at the smart-contract layer can be nullified by semiconductor bottlenecks below it or statutory compliance mismatches above it.
Main Takeaways:

- **Hardware & IP Chokepoints (L0 ↔ L2):** Custom FHE ASICs rely on concentrated sub-5nm foundries, optical chips face packaging and conversion overhead, and proprietary patent walls separate enterprise PET vendors from open Web3 protocols.
- **Metadata & Swarm Containment (L2 ↔ L2.5):** Encrypting invoice amounts is insufficient if on-chain transaction graphs leak business relationships, while TEE-based AI swarms must be constrained by deterministic mathematical invariants so model hallucinations cannot alter financial balances.
- **Regulatory Compliance Gateways (L2.5 ↔ L3):** Under EU PSD3/PSR, DORA, and national set-off laws (such as Portuguese DL 150/2019), financial institutions cannot execute automated smart-contract clearing without licensed PISP/AISP bridges, auditable ICT SLAs, and revocable cryptographic access controls.

5 Ecosystem Verticalizations and Layer 3 Institutional Hybridization

To test Hypothesis H_3 logically, we examine how industrial specialization and statutory pressures are reshaping the boundary between decentralized cryptographic protocols (Layers 2 and 2.5) and incumbent enterprise and banking institutions (Layer 3). Rather than wholesale replacement of legacy financial rails by public blockchains, empirical evidence indicates an accelerating trend toward *vertical integration* and *institutional hybridization* [1, 9].

5.1 Verticalization Vectors Across the PET Stack

Because cross-layer friction (ISA fragmentation at Layer 0, IP licensing cliffs at Layer 1, and DORA/PSD3 compliance at Layer 3) penalizes loosely coupled architectures, leading ecosystem actors are consolidating along three primary verticalization vectors:

1. **Full-Stack Sovereign Liquidity & Clearing Verticals:** Exemplified by the *FluidCluster / Ecompensa* architecture, this vertical couples Layer 0 silicon-photonic MZI acceleration directly with Layer 2 `fhEVM` optimization engines, Layer 2.5 TEE-attested MCP agent swarms, and Layer 3 regulatory sandboxes (e.g., Banco de Portugal and Iberian e-invoicing networks) [1, 4]. By owning the end-to-end cryptographic pipeline from local ERP ingestion to residual PISP settlement, full-stack verticals eliminate third-party middleware latency and satisfy DORA ICT concentration rules under a single auditable governance umbrella.

2. **Institutional DeFi & Confidential Stablecoin Verticals:** Bridged by protocols such as Zama, Fhenix, and Inco alongside compliant stablecoin issuers, this vertical wraps regulated fiat-backed stablecoins into `ERC-7984` confidential ciphertokens (`cUSDC`, `cUSDT`) [8, 36]. By embedding programmable zero-knowledge AML/KYB gates and selective regulatory viewing keys (`FHE.allow`) directly into the token contract, institutional treasuries gain access to dark-pool liquidity and atomic delivery-versus-payment (DvP) settlement without exposing corporate treasury positions on public blockchains.

3. **Multi-Party Confidential AI & Corporate M&A Verticals:** Commercialized by enterprise platforms (Duality Technologies, Opaque) and specialized Web3 protocols (Fluid-

Cluster M&A Aperture), this vertical integrates federated encrypted analytics with game-theoretic mechanism design [26, 7]. Competing enterprises jointly train supply-chain demand models or evaluate cross-company merger synergies (computing Shapley value equilibria over encrypted balance sheets) without revealing proprietary customer lists or margin structures prior to binding closing [6].

5.2 Hybrid Archetype I: The Collaborative FHE-Based ERP

Traditional Enterprise Resource Planning (ERP) suites (SAP S/4HANA, Oracle NetSuite, Microsoft Dynamics 365, Primavera/PHC) operate as isolated data fortresses. While internal processes are automated, cross-enterprise coordination relies on insecure plaintext EDI/PDF exchanges or centralized supply-chain portals that expose sensitive pricing and inventory buffers—directly fueling the Bullwhip Effect [10, 9].

Hybridizing Layer 3 ERP systems with Layer 2 `fhEVM` and Layer 2.5 agentic swarms produces the **Collaborative FHE-Based ERP**: a zero-trust cooperative overlay that preserves existing on-premise or cloud ERP databases as local systems of record while enabling global encrypted execution across 10 core inter-enterprise modules [9]:

- **Cryptographic Sidecar Ingestion:** A lightweight local ERP connector reads accounts payable, accounts receivable, and inventory telemetry inside the enterprise firewall. Instead of transmitting raw invoices, the sidecar hashes counterparty identifiers for off-chain Private Set Intersection (PSI) [5], encrypts monetary values under the network FHE public key (`ERC-7984` ciphertokens), and generates a client-side ZK-SNARK attesting to tax-authority validity (e.g., SAF-T PT digital signature) [6].
- **Blind Multilateral Netting & Bullwhip Dampening:** Once off-chain PSI detects closed debt cycles ($A \to B \to C \to A$), the Layer 2 `fhEVM` engine executes homomorphic Leontief / Network Simplex optimization [2, 4]. Simultaneously, upstream and downstream suppliers feed encrypted order-book velocity into a homomorphic Bullwhip incentive solver, aligning production schedules without revealing unit margins or end-customer identities [10].
- **Automated Cooperative Modules:** Beyond liquidity netting, the Collaborative FHE-ERP supports blind joint procurement (aggregating encrypted demand across SMEs to unlock Tier-1 volume discounts), encrypted ESG/Scope-3 carbon accounting, and automated debt-to-equity restructuring proposals mediated by Layer 2.5 MCP AI swarms [9, 38].
- **Atomic Ledger Reconciliation:** Upon threshold decryption via the MPC-KMS committee, each participating ERP automatically posts cryptographically signed accounting journal entries (complying with statutory compensation rules such as Portuguese Decree-Law no. 150/2019), reconciling settled obligations without human intervention.

5.3 Hybrid Archetype II: The Photonic-Mutual Bank (FHEVM Sovereign Mutual)

While the Collaborative FHE-ERP transforms how enterprises coordinate commercial obligations, final settlement of net residual balances still depends on banking intermediaries. Traditional fractional-reserve and clearing banks operate as centralized plaintext custodians: they observe every counterparty payment, monetize proprietary cash-flow telemetry, and impose high working-capital fees on Small and Medium Enterprises (SMEs).

Synthesizing a cooperative credit mutual with Layer 0 opto-photonic acceleration, Layer 2 `fhEVM` smart contracts, and Layer 3 PSD3/PSR open-banking licenses yields a novel institutional archetype: the **Photonic-Mutual Bank** [1, 8, 40]. Operating under the principle of *Absolute*

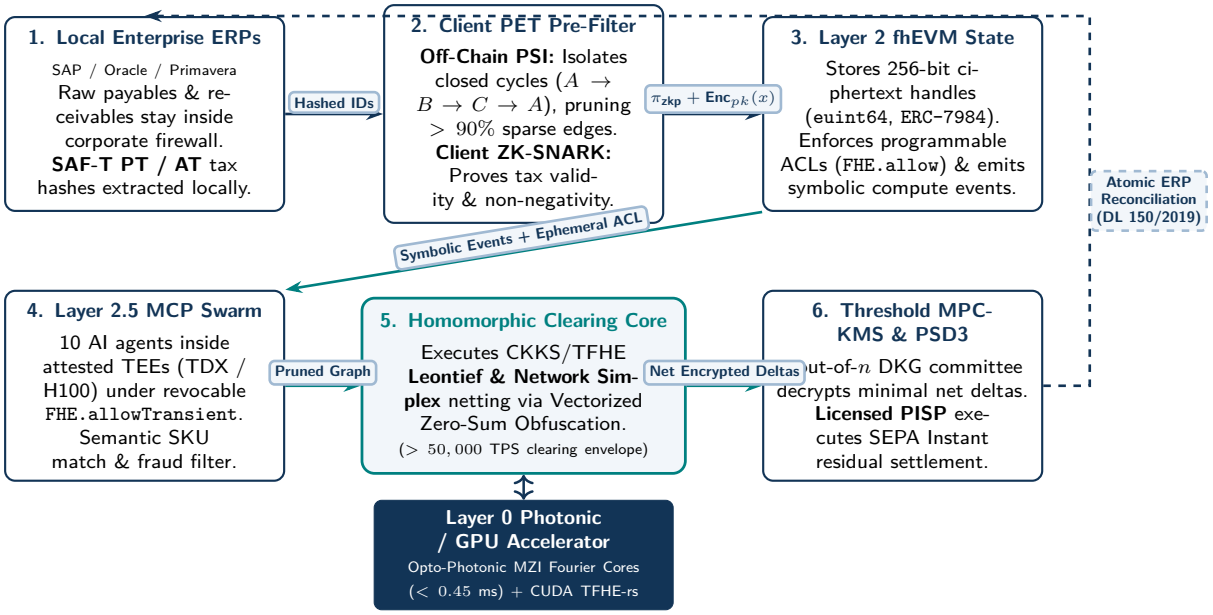


Figure 5: FluidCluster Blind Triangle & Off-Chain Coprocessor Execution Loop: decoupling off-chain PSI topology discovery and client-side ZK-SNARK validation from `fhEVM` consensus, photonic-accelerated Simplex netting, and Threshold MPC-KMS / PSD3 settlement.

Computational Blindness, the Photonic-Mutual Bank separates fiduciary liquidity stewardship from plaintext data custody:

1. **Cryptographically Blind Balance Sheet & Mutual Liquidity Pooling:** Corporate members deposit fiat liquidity into segregated safeguard accounts at partner Account Servicing Payment Service Providers (ASPSPs) or central bank settlement facilities, minting 1:1 backed confidential ciphertokens (`ERC-7984` / `cEUR`) [36, 8]. Member balances, credit lines, and bilateral trade exposures remain continuously encrypted on-chain under RLWE lattice cryptography.

2. **Sub-Millisecond Opto-Photonic Clearing Core:** Rather than batching overnight ACH/SEPA clearing files, the mutual operates a Layer 0 opto-photonic coprocessor cluster (utilizing Mach-Zehnder Interferometer optical Fourier cores) coupled to its `fhEVM` state engine [4, 22]. Continuous multilateral netting and homomorphic credit-risk scoring execute in < 0.45 ms per batch at $> 50,000$ TPS with near-zero thermal energy overhead, recycling internal member liquidity before external interbank rails are ever invoked [1, 2].

3. **Threshold MPC-KMS Governance & Supervisory Auditability:** Governance follows an `fheDAO` cooperative charter where member enterprises vote homomorphically on reserve ratios and credit-pool parameters [37]. Decryption keys are partitioned across a multi-institutional Threshold MPC-KMS committee comprising elected member representatives, independent auditors, and national supervisory observers (e.g., Banco de Portugal / ECB) [11, 3]. Regulators hold programmable `FHE.allow` audit keys to verify capital adequacy, liquidity coverage ratios (LCR), and AML invariants in real time without exposing individual SME trade secrets to commercial competitors.

4. **Native PSD3 PISP/AISP Settlement Bridge:** When multilateral netting leaves a minimal non-zero residual debt, the Photonic-Mutual Bank's licensed Layer 2.5/3 Payment Initiation Service Provider (PISP) module triggers an automated, Strong Customer Authentication (SCA)-compliant SEPA Instant transfer directly between member ASPSP accounts [29, 30].

5.4 Migration Roadmap for Incumbent Banks, PSPs, and ERP Vendors

Because incumbent Layer 3 institutions cannot replace core banking ledgers or mission-critical ERP databases overnight without violating DORA operational continuity mandates, transition toward FHE-native hybrid architectures follows a three-phase incremental migration roadmap:

Phase 1: TEE-Enclave Sidecars & Read-Only AISP Telemetry (Months 1–9):
Incumbent banks and ERP vendors deploy hardware-attested TEE enclaves (Intel TDX / NVIDIA Confidential Compute) alongside existing databases, ingesting real-time cash-flow telemetry via PSD3 Account Information Service Provider (AISP) APIs [28, 29]. Off-chain Private Set Intersection (PSI) identifies multilateral debt cycles across participating corporate clients without modifying core ERP schemas [5].

Phase 2: Hybrid fhEVM Netting Overlay & PISP Residual Settlement (Months 9–24):
Encrypted cycle subgraphs are anchored onto a permissioned or rollup-based `fhEVM` settlement layer using `ERC-7984` ciphertokens and client-side ZK-SNARK tax proofs (SAF-T / EN 16931) [6, 9]. Threshold MPC-KMS governance is established across consortium banks, and residual net balances are settled automatically via licensed PISP payment initiation triggers [11, 30].

Phase 3: Photonic-Accelerated FHE Core Banking & Autonomous Swarm Orchestration (Months 24+):
High-volume clearing houses and cooperative mutuals integrate Layer 0 opto-photonic PICs and custom FHE ASICs to eliminate GPU bootstrapping bottlenecks [1, 21]. Full-spectrum cooperative ERP modules—including blind supply-chain Bullwhip stabilization, automated M&A coalition netting, and continuous regulatory zero-knowledge reporting—operate natively over encrypted state [10, 7].

Chapter 5 Explainer: The Future Hybrid Enterprise & Bank

Core Thesis: Layer 3 institutions will not be replaced by public blockchains; instead, they will absorb Layers 2 and 2.5 to become cryptographically blind mutual networks that unlock trapped working capital without surrendering commercial privacy.
Main Takeaways:

- **Collaborative FHE-ERP:** Enterprises keep their raw data inside existing ERP systems (SAP, Oracle, Primavera) while using cryptographic sidecars to net mutual debts, coordinate supply chains against the Bullwhip Effect, and negotiate joint procurement without exposing prices or suppliers.
- **Photonic-Mutual Bank:** Combining optical FHE accelerators (Layer 0) with `fhEVM` smart contracts (Layer 2) and PSD3 open-banking licenses (Layer 3) enables a cooperative bank that clears trade debt in under 0.45 ms while remaining mathematically blind to customer trade secrets.
- **Phased Institutional Adoption:** Banks and ERP vendors can migrate safely in three stages—starting with TEE/AISP read-only sidecars, advancing to hybrid `fhEVM` + PISP netting, and culminating in photonic-accelerated FHE core infrastructure.

6 Synthesis, Risk Scoring Matrix, and Future Research Agenda

6.1 Hypothesis Validation Summary (H_1, H_2, H_3)

Synthesizing our architectural decomposition across Layers 0 through 3 and our STRIDE-PET supply-chain risk evaluation yields structured logical validation across all three foundational hypotheses formulated in Section 1.3:

1. **Validation of Hypothesis H_1 (The Trust-Performance Inversion): *Confirmed.*** Comparative benchmarking demonstrates an inverse structural trade-off between hardware-rooted TEEs and lattice-based FHE. Pure TEE architectures (Intel TDX, AMD SEV-SNP, NVIDIA Confidential Compute) achieve sub-millisecond clearing latency ($\tau_{\text{clear}} < 1$ ms) with minimal ciphertext bloat ($\Omega_{\text{BW}} \approx 1\times$), but concentrate systemic risk inside an irreducible Trusted Computing Base ($\mathcal{S}_{\text{TCB}}$) vulnerable to microarchitectural side-channels and single-vendor foundry control [19, 27]. Conversely, pure FHE eliminates hardware trust assumptions via post-quantum RLWE hardness, yet incurs 10^3–$10^5\times$ wire expansion and heavy bootstrapping overhead that remain economically unviable on general-purpose CPUs [14]. Offloading polynomial NTT/FFT transforms to specialized Layer 0 custom ASICs (Niobium/SEMIFIVE, Cornami) and opto-photonic MZI cores (Optalysys, FluidCluster PIC solvers) compresses core optimization latency to < 0.45 ms, resolving the inversion without restoring hardware plaintext trust [1, 4, 22].

2. **Validation of Hypothesis H_2 (Modular Hybrid Convergence): *Confirmed.*** Architectural evaluation of the FluidCluster 2.0/3.0 pipeline confirms that composing complementary PETs across well-defined cryptographic isolation boundaries outperforms any monolithic paradigm [6, 19]. Executing multi-party Private Set Intersection (PSI) off-chain prunes sparse enterprise invoice graphs by $> 90\%$ prior to on-chain submission, simultaneously avoiding $O(N^2)$ FHE gate blowup and concealing non-intersecting commercial relationships from graph-topology inference [5, 15]. Binding client-side ZK-SNARKs to `ERC-7984` ciphertext commitments guarantees tax-authority invoice authenticity (SAF-T/AT) before homomorphic evaluation begins, while confining Layer 2.5 MCP AI swarms inside attested TEEs under revocable `FHE.allowTransient` grants enables semantic anomaly filtering without exposing the master Threshold MPC-KMS decryption key [35, 12].

3. **Validation of Hypothesis H_3 (Layer 3 Institutional Hybridization): *Logically Substantiated.*** Mapping EU PSD3/PSR open-banking rules, DORA ICT third-party risk mandates, and national set-off statutes (Portuguese Decree-Law no. 150/2019) onto Layer 2.5/3 interfaces confirms that public, unpermissioned blockchains cannot directly substitute for regulated banking rails, nor can centralized plaintext clearinghouses satisfy modern trade-secret and operational resilience demands [29, 30, 10]. Institutional convergence around two hybrid archetypes—the *Collaborative FHE-Based ERP* and the *Photonic-Mutual Bank*—provides a legally compliant migration pathway that preserves existing ERP systems of record and licensed PISP/AISP settlement endpoints while shifting multilateral liquidity clearing into a cryptographically blind domain [9, 8].

6.2 Composite Supply Chain Risk Scorecard Across Architectures

To operationalize our Socio-Technical Supply Chain Risk Management (STS-SCRM) findings for enterprise architects, central bank sandboxes, and institutional risk committees, Table 2 contrasts composite risk exposure across three deployment archetypes: (A) Pure Cloud TEE Clearing, (B) Pure Unaccelerated On-Chain FHE, and (C) the FluidCluster Hybrid Architecture. Risk severity is scored on a standardized five-point ordinal scale (**Low**, **Low–Moderate**, **Moderate**, **High**, **Critical**) across seven cross-layer risk vectors.

Table 2: Composite Supply Chain & STRIDE-PET Risk Scorecard Across Enterprise Clearing Architectures

Supply Chain & Threat Vector	**Archetype A: Pure Cloud TEE Clearing** (Intel TDX / NVIDIA H100 Enclaves)	**Archetype B: Pure On-Chain FHE** (Unpruned CPU/GPU fhEVM)	**Archetype C: FluidCluster Hybrid** (PSI + ZKP + fhEVM + MCP + PIC)
1. Silicon Foundry & Hardware Trust Risk	**Critical** — Total reliance on CPU/GPU vendor fuses, microcode updates, and East Asian sub-5nm foundries.	**Low** — Cryptographic security rests on RLWE math; hardware untrusted for confidentiality.	**Low–Moderate** — Confidentiality anchored in RLWE math; Layer 0 photonic/GPU redundancy mitigates liveness risk.
2. Side-Channel & Enclave Extraction Risk	**High** — Exposed to cache-timing, voltage fault injection, and co-resident hypervisor leaks.	**Low** — No plaintext ever materialized in processor registers or memory.	**Low** — Enclaves restricted to ephemeral Layer 2.5 MCP swarm micro-contexts with revocable ACLs.
3. Computational Bottleneck & Scalability Risk	**Low** — Near-native execution speeds ($<$ 1 ms) across large graphs.	**Critical** — $O(N^2)$ ciphertext bloat and slow bootstrapping stall throughput ($<$ 10 TPS).	**Low–Moderate** — Off-chain PSI prunes $>$ 90% of edges; photonic MZI solver clears batches in $<$ 0.45 ms.
4. Topological Metadata & Graph Leakage	**Moderate–High** — Host OS and network telemetry reveal counterparty communication patterns.	**High** — On-chain ciphertext handles and `FHE.allow` events leak business graph topology.	**Low** — Off-chain PSI isolates cycles; batch zero-sum netting masks individual edge weights.
5. Key Custody & Insider Collusion Risk	**High** — Single KMS or cloud enclave operator compromise exposes historical plaintext state.	**Low–Moderate** — Protected if governed by multi-party Threshold MPC-KMS.	**Low** — Multi-jurisdictional t-out-of-n Threshold MPC-KMS with proactive share rotation and slashing.
6. IP Licensing & Compiler Lock-In (H_{lock})	**Moderate** — Tied to proprietary vendor attestation SDKs (TDX/SEV-SNP/Nitro).	**Low–Moderate** — Open-source stacks (`TFHE-rs`, `OpenFHE`) reduce software lock-in.	**Moderate** — Open-source `fhEVM`/ZKP core balanced against specialized Layer 0 photonic/ASIC interfaces.
7. Regulatory Compliance (PSD3 / DORA / GDPR)	**Moderate** — Meets DORA cloud SLAs but lacks mathematical verifiability for external auditors.	**High** — Permissionless nodes clash with DORA ICT rules, PSD3 SCA, and GDPR erasure.	**Low** — Permissioned co-processor SLAs, licensed PISP/AISP bridges, ZK-AML, and cryptographic erasure.

6.3 Open Problems and Next-Generation Research Directions

While the FluidCluster hybrid architecture resolves the immediate privacy-latency-verifiability trilemma for multilateral clearing and collaborative supply-chain optimization, our cross-layer supply-chain risk analysis exposes four fundamental research frontiers that require coordinated academic, industrial, and regulatory progress:

1. **Succinct Verifiable FHE (`vFHE`) for Trustless Coprocessor Execution:** In current symbolic `fhEVM` architectures, smart contracts rely on economic slashing, multi-operator replication, or TEE attestations to trust that off-chain coprocessors executed TFHE/CKKS bootstrapping honestly [13, 16]. Generating recursive Zero-Knowledge Proofs (SNARKs/STARKs) directly over lattice bootstrapping circuits—known as *Verifiable FHE (vFHE)*—remains computationally expensive due to field-arithmetic mismatches between RLWE cyclotomic rings and elliptic-curve/hash proving systems. Advancing hardware-accelerated vFHE on unified cryptographic chips (such as Fabric's VPU or photonic-electronic hybrid cores) is essential to eliminate coprocessor trust assumptions entirely.

2. **Cross-Foundry Intermediate Representation (IR) for Optical-Electronic FHE:** To eliminate Layer 0 vendor lock-in (H_{lock}) across GPUs, FPGAs, custom FHE ASICs (Niobium, Cornami, Chain Reaction), and silicon-photonic MZI accelerators (Optalysys), the ecosystem requires a standardized, hardware-agnostic Intermediate Representation (extending MLIR / HEIR and FHETCH consortium specifications [21]). Such an IR must automatically partition arithmetic circuits between electronic logic (for non-linear TFHE programmable bootstrapping) and optical Fourier cores (for high-degree linear polynomial convolutions) while calibrating analog-to-digital precision noise budgets dynamically.

3. **Post-Quantum End-to-End Composability Across Hybrid Stacks:** Although TFHE and CKKS provide native post-quantum security under Ring Learning With Errors (RLWE) assumptions, hybrid pipelines frequently couple FHE ciphertexts with pre-quantum elliptic-curve ZK-SNARKs (e.g., Groth16/Plonk over BN254) and classical ECDSA TEE attestation roots [19]. Transitioning client-side invoice validity proofs and Threshold MPC-KMS key-share protocols to fully post-quantum lattice or hash-based primitives (such as Lattice-SNARKs and ML-DSA/Dilithium threshold signatures) is mandatory to prevent "harvest-now, decrypt-later" attacks against long-lived corporate trade secrets.

4. **Automated Regulatory Zero-Knowledge Compliance Circuits (PSD3 / MiCA / DORA):** As European financial regulation evolves under PSD3/PSR, DORA, and Markets in Crypto-Assets (MiCA), future research must formalize statutory legal rules into machine-verifiable zero-knowledge compliance circuits [29, 8]. Embedding standardized ZK-AML/CFT proofs, real-time Liquidity Coverage Ratio (LCR) attestations, and automated SAF-T tax-audit triggers directly into `ERC-7984` confidential ciphertokens will enable central banks and supervisory authorities to exercise continuous macroprudential oversight without compromising enterprise micro-privacy.

Chapter 6 Explainer: Executive Conclusion & Strategic Roadmap

Core Thesis: Hybrid Privacy-Enhancing Technology (PET) architectures—governed by threshold cryptography and accelerated by specialized photonic and ASIC hardware—represent the optimal socio-technical equilibrium for sovereign enterprise and financial clearing infrastructure.

Main Takeaways:

- **No Single-Primitive Silver Bullet:** Pure cloud TEEs concentrate catastrophic risk in silicon vendors and side-channel surfaces, while unaccelerated pure FHE collapses

under ciphertext bloat and metadata leakage.

- **Validated Separation of Duties:** FluidCluster's layered architecture proves that combining off-chain PSI (for cycle pruning), client-side ZKPs (for tax/invoice validity), `fhEVM` + Photonic MZI solvers (for blind optimization), Threshold MPC-KMS (for distributed key governance), and TEE-isolated MCP swarms (for semantic normalization) achieves sub-millisecond clearing with institutional-grade privacy.
- **The Institutional Horizon:** As Collaborative FHE-ERPs and Photonic-Mutual Banks mature under PSD3/PSR and DORA, confidential computing will shift from a defensive cybersecurity add-on to the foundational operating system of global B2B liquidity and supply-chain coordination.

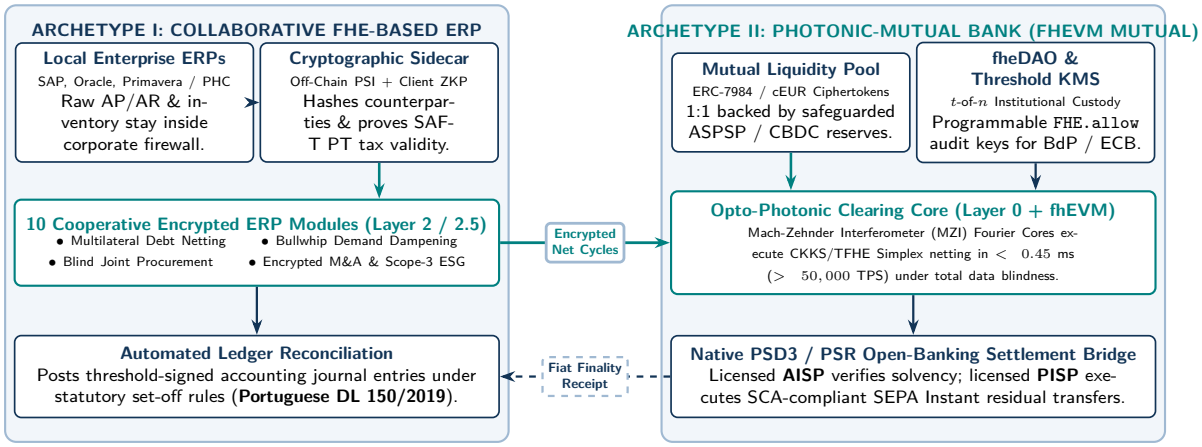


Figure 6: Layer 3 Institutional Hybridization: architectural coupling between the *Collaborative FHE-Based ERP* (left) and the *Photonic-Mutual Bank* (right), uniting local ERP privacy, photonic-accelerated `fhEVM` clearing, Threshold MPC-KMS regulatory auditability, and PSD3 PISP/AISP settlement.

References

[1] FluidCluster Research Group, "FluidCluster 2.0: The Sovereign Operating System for European Liquidity — Integrating TFHE/CKKS, Photonic Integrated Circuits, and MCP AI Swarms," *FluidCluster Technical Whitepaper Series*, Oct. 2026. `https://www.fluidcluster.pt/`

[2] FluidCluster Research Group, "Hardware-Accelerated FHE Netting Engine via Vectorized Zero-Sum Obfuscation and Photonic Coprocessing," *FluidCluster Technical Report*, Sept. 2026. `https://fluidcluster.pt/assets/docs/fluidcluster_netting_paper.pdf`

[3] FluidCluster Research Group, "Mathematical and Game-Theoretic Foundations of Sovereign Liquidity Clearing," *FluidCluster Whitepaper*, 2026. `https://fluidcluster.pt/assets/docs/sovereign-liquidity-clearing-whitepaper.pdf`

[4] FluidCluster Research Group, "Generalizing Linear Programming Minimum-Cost Flow Problems for Homomorphic Silicon-Photonic Solvers," *FluidCluster Computational Physics Optimization Series*, 2026.

[5] FluidCluster Research Group, "Accelerating fhEVM With Off-Chain Private Set Intersection (PSI): System Constraints Under Load," *FluidCluster Telemetry Architecture Paper*, Aug. 2026.

[6] FluidCluster Research Group, "FluidCluster 3.0 MA Aperture Protocol: Unifying Leontief Network Simplex Solvers With fhEVM Smart Contracts and Client-Side ZK-SNARKs," *FluidCluster Sovereign DLT Series*, Oct. 2026.

[7] FluidCluster Research Group, "Automated Confidential MA Integration: Mechanism Design Stable Coalition Incentives via fhEVM and Shapley Value Equilibria," *FluidCluster Academic Series*, Sept. 2026.

[8] FluidCluster Research Group, "Strategic Report: Confidential Stablecoins (cUSDC cUSDT) and Sovereign Clearing Infrastructure," *FluidCluster Market Technology Intelligence*, Sept. 2026. `https://fluidcluster.pt/assets/docs/Confidential_Stablecoins_Analysis_Report.pdf`

[9] FluidCluster Research Group, "FluidCluster ERP: The 10 Core Cooperative Modules Disrupting Siloed Enterprise Systems via Confidential Deep Tech PETs," *FluidCluster Enterprise Architecture Brief*, Nov. 2026.

[10] FluidCluster Research Group, "Taming the Bullwhip Effect via FHE: Blind Supply Chain Coordination and Incentive Solvers (DL 150/2019 DORA Compliance)," *FluidCluster Whitepaper*, Sept. 2026.

[11] Zama, "fhEVM: A Full-Stack Framework for Integrating Fully Homomorphic Encryption with EVM-Compatible Blockchains," *Zama Protocol Documentation GitHub*, 2026. `github.com`

[12] Zama, "FHEVM Solidity Guides: Encrypted Types, Symbolic Execution Coprocessors, and ACL Delegation," *Zama Docs*, May 2026. `zama.org`

[13] Blockworks Research, "Understanding Zama: fhEVM Coprocessor Architecture and Modular ZK Integration (Kakarot Acquisition)," *Blockworks Intel Report*, Mar. 2026.

[14] Zama, "TFHE Bootstrapping in Under 1 Millisecond on GPU," *Zama Engineering Blog*, 2024. `zama.org`

[15] A. Author et al., "Can Fully Homomorphic Encryption on Public Blockchains Preserve Application Privacy? Access Control Delegation and Metadata Leakage in fhEVM," *IACR Cryptology ePrint Archive*, Report 2026/2771, Nov. 2026. `iacr.org`

[16] S. Author et al., "SoK: Fully Homomorphic Encryption in Smart Contract Environments and Economic Mechanism Design," *IACR Cryptology ePrint Archive*, Report 2025/527, Mar. 2025. `iacr.org`

[17] Fhenix, "Fhenix Acquires Sunscreen to Advance Quantum-Safe FHE and ZKP Compiler Infrastructure on Ethereum, Arbitrum, and Base," *Fhenix Research Announcement*, July 2026. `fhenix.io`

[18] Inco Network, "Confidential Compute Server Architecture: Combining TEEs and fhEVM for Verifiable Blockchain Privacy," *Inco Protocol Documentation*, 2026. `https://docs.inco.org/architecture/components`

[19] Aztec Research, "Is ZK-MPC-FHE-TEE a Real Creature? Composability and Trade-Offs of Multi-PET Architectures," *Aztec Network Technical Blog*, Mar. 2025. `aztec.network`

[20] Phala Network, "TEE vs. MPC vs. ZK: Architectural Complementarity in Confidential Computing," *Phala Research*, Apr. 2025. phala.com

[21] Optalysys, Chain Reaction, and Niobium Microsystems, "Launch of the FHE Technical Consortium for Hardware (FHETCH) for Cross-Stack Interoperability," Oct. 2024. optalysys.com

[22] Optalysys, "An Optical Hardware Accelerator for Lattice-Based Cryptography and FHE via Optical Fourier Transforms," *FHE.org Technical Meetup*, 2023. fhe.org

[23] SEMIFIVE Niobium Microsystems, "SEMIFIVE Partners with Niobium to Develop Custom FHE Hardware Accelerator ASIC," *Industry Press Release*, Feb. 2026.

[24] V. Chadha, "Full Circuit FHE Acceleration on Cornami Scalable Non-Von Neumann Fabric (SCIFR API)," *ITU-T Workshop on Quantum-Safe Confidential Computing*, 2022.

[25] Communications of the ACM, "Homomorphic Technologies Could Process Still-Encrypted Data: Survey of Cornami, Intel, Niobium, and Optalysys," *CACM News*, June 2025.

[26] Duality Technologies, "Secure Data Collaboration Platform: OpenFHE, Federated Learning, and Encrypted Analytics Across Trust Boundaries," 2026. dualitytech.com

[27] Fortanix, "Confidential Computing Technology and Hardware-Enforced TEE Orchestration," 2026. https://www.fortanix.com/faq/confidential-computing/confidential-computing-technology

[28] Decryption Digest, "Fortanix vs. Opaque vs. Anjuna: Enterprise Confidential Computing and GPU TEE Enablement Compared," Sept. 2026.

[29] Herbert Smith Freehills Kramer, "The New EU Payments Rulebook (PSD3/PSR): PISP/AISP Open Banking Access and Prohibited Obstacles," *FSR Corporate Crime Notes*, May 2026.

[30] Compliance Support NL, "PSD3 PSR Operational Impact on PISPs, AISPs, and Account-Servicing Banks," Nov. 2026.

[31] T. J. Pettit, K. L. Croxton, and J. Fiksel, "The Evolution of Resilience in Supply Chain Management: A Retrospective on Ensuring Supply Chain Resilience," *Journal of Business Logistics*, vol. 40, no. 1, pp. 56–65, 2019.

[32] A. Wicaksana et al., "Supplier Risk in Supply Chain Risk Management: An Updated Conceptual Framework," *Applied Sciences*, June 2025.

[33] A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design Science in Information Systems Research," *MIS Quarterly*, vol. 28, no. 1, pp. 75–105, Mar. 2004.

[34] FluidCluster Research Group, "FluidCluster Institutional Pitch & Regulatory Sandbox Architecture (DORA & Banco de Portugal Telemetry)," *FluidCluster Sovereign Liquidity Series*, Oct. 2026. https://www.fluidcluster.pt/legal.html

[35] FluidCluster Research Group, "Cognitive MCP Swarms in Hardware-Attested TEEs: Autonomous Invoice Normalization and Anomaly Filtering for fhEVM Clearing," *FluidCluster Technical Report*, Oct. 2026. https://www.fluidcluster.pt/apps.html

[36] FluidCluster Research Group, "Confidential RWA Tokenization via ERC-7984: Cryptographic Receivable Ciphertokens and Privacy-Preserving Settlement," *FluidCluster Technical Brief*, Sept. 2026. https://www.fluidcluster.pt/apps.html

[37] FluidCluster Research Group, "Blind Sovereign Governance: Homomorphic Cooperative Charters and Anti-Coercion Voting via fheDAOs," *FluidCluster Governance Series*, Oct. 2026.

[38] FluidCluster Research Group, "FluidCluster Fusion: Autonomous Debt-to-Equity Swaps and Balance-Sheet Purification Under Portuguese Decree-Law no. 150/2019," *FluidCluster Whitepaper Series*, Oct. 2026. `https://www.fluidcluster.pt/fusion.html`

[39] Niobium Microsystems, "Hardware-Accelerated Fully Homomorphic Encryption (DARPA DPRIVE): Custom Silicon for Zero-Trust Cloud Compute," *Niobium Technical Whitepaper*, 2026.

[40] Zama Research, "Confidential Banking and Institutional Settlement on fhEVM: Encrypted Balances, Regulatory Viewing Keys, and Threshold KMS," *Zama Protocol Architecture Series*, 2026.