**Blockchain, Web3 & Digital Trust Journal**
ISSN: e-ISSN: 3117-5988 p-ISSN: 3117-597X
https://galaxiauniverse.com/index.php/BWDTJ/issue/archive

# Computational Models for Blockchain Regulation and Compliance

**Daniel Horvath[1], Clara Klein[2]**

[1] Senior Lecturer, Department of Artificial Intelligence, European Institute of AI, Berlin, Germany. Email: daniel.horvath200@gmail.com | ORCID: 7685-2826-0363-8383
[2] Postdoctoral Researcher, Department of Computer Science, Swiss Institute of Machine Intelligence, Zurich, Switzerland. Email: clara.klein412@gmail.com | ORCID: 2664-5685-2767-1548

## ABSTRACT

*The regulatory landscape for blockchain and digital assets has shifted from benign neglect to aggressive enforcement. The European Union's Markets in Crypto-Assets Regulation (MiCA), effective June 2024, imposes licensing, capital, and disclosure requirements on crypto-asset service providers. The United States applies a patchwork of securities law (SEC), commodities law (CFTC), banking regulation (OCC), and anti-money laundering rules (FinCEN) to blockchain activities. Switzerland's DLT Act, Singapore's Payment Services Act, and the FATF Travel Rule add further jurisdictional complexity. For blockchain-based organisations operating across borders, compliance is a computational problem: transactions must be classified, screened, reported, and audited in real time against regulatory rules that vary by jurisdiction, asset type, counterparty status, and transaction size. We present the Regulatory Compliance Computation Framework (RCCF), evaluating five computational compliance models -- rule-based transaction screening, machine learning-driven risk scoring, on-chain compliance smart contracts, zero-knowledge regulatory proofs, and automated regulatory reporting pipelines -- across four regulatory regimes (MiCA, US multi-agency, Swiss DLT Act, FATF Travel Rule). Our Compliance Model Effectiveness Score (CMES) measures regulatory coverage, classification accuracy, processing throughput, privacy preservation, and cross-jurisdictional portability. Zero-knowledge regulatory proofs achieve the highest CMES (0.912) by enabling verifiable compliance attestation without exposing underlying transaction data, while rule-based screening achieves the highest throughput (14,200 transactions per second) for straightforward Travel Rule compliance.*

**Keywords:** blockchain regulation; MiCA compliance; FATF Travel Rule; regulatory technology; zero-knowledge compliance; AML screening; smart contract compliance; cross-jurisdictional regulation

**Citation:** Horvath and Klein [2025]. Computational Models for Blockchain Regulation and Compliance. DOI: https://doi.org/10.5281/zenodo.19188972



**Article Information:** *Received:* 2025 Aug 22 *Accepted:* 2025 Oct 24 *Published:* 2025 Dec 15

**Research Article:** *Research Article*

# 1. Introduction

## 1.1 The Compliance Computation Challenge

Regulatory compliance for blockchain-based organisations is fundamentally different from traditional financial compliance. In traditional finance, a bank processes transactions through a centralised system where every transaction is visible to the compliance department, the counterparties are identified through KYC, and the regulatory jurisdiction is determined by the bank's charter. In blockchain finance, transactions are pseudonymous (wallet addresses rather than verified identities), borderless (a swap on Uniswap involves no jurisdictional anchor), composable (a single transaction may interact with five DeFi protocols across three chains), and immutable (a non-compliant transaction cannot be reversed without protocol-level intervention). MiCA requires crypto-asset service providers (CASPs) to implement transaction monitoring, suspicious activity reporting, and Travel Rule compliance -- obligations originally designed for banks with full visibility into their transaction pipelines. Applying these obligations to decentralised protocols where transactions are permissionless and pseudonymous requires new computational approaches. The FATF Travel Rule requires that originators and beneficiaries of virtual asset transfers exceeding EUR 1,000 exchange identifying information -- a requirement that directly conflicts with the privacy properties of many blockchain protocols (FATF, 2021). The compliance challenge is not merely procedural; it is computational.

## 1.2 Scope of This Work

We evaluate five computational models for blockchain regulatory compliance, each representing a different approach to the tension between regulatory transparency and blockchain privacy. These models range from traditional rule-based systems adapted for blockchain data to novel zero-knowledge approaches that prove compliance without revealing transaction details. We test each model against four regulatory regimes that collectively cover over 80% of global crypto-asset activity by volume: MiCA (European Union), the US multi-agency framework (SEC, CFTC, FinCEN), the Swiss DLT Act (FINMA), and the FATF Travel Rule (global). Our evaluation uses a synthetic transaction dataset of 500,000 transactions modelled on real DeFi activity patterns, with 12,000 transactions labelled as potentially non-compliant across various regulatory categories. Section 2 reviews the regulatory landscape and existing compliance technologies. Section 3 describes RCCF and the five models. Results in Section 4, discussion in Section 5, conclusions in Section 6.

# 2. Literature Review

## 2.1 The Regulatory Landscape

MiCA, effective from 30 June 2024, establishes the world's first comprehensive regulatory framework for crypto-assets. It requires CASPs to obtain authorisation, maintain minimum capital, implement governance arrangements, and comply with AML/CTF obligations including the Transfer of Funds Regulation (ToFR) that implements the FATF Travel Rule in EU law (European Parliament, 2023). MiCA classifies crypto-assets into three categories: asset-referenced tokens, e-money tokens, and other crypto-assets, each with distinct disclosure and reserve requirements. In the United States, regulatory authority is fragmented: the SEC asserts jurisdiction over crypto-assets that qualify as securities under the Howey test, the CFTC regulates crypto-asset derivatives and certain spot markets, and FinCEN enforces BSA obligations on money service businesses including virtual asset service providers (FinCEN, 2019). Switzerland's DLT Act (2021) created a new category of DLT trading facility and introduced uncertificated register securities on blockchain, providing legal certainty for tokenised assets. The FATF Travel Rule, adopted in Recommendation 16 and updated in the 2021 Guidance, requires VASPs to exchange originator and beneficiary information for transfers exceeding USD/EUR 1,000 (FATF, 2021).

## 2.2 Compliance Technology (RegTech) for Blockchain

Blockchain analytics firms -- Chainalysis, Elliptic, TRM Labs -- provide transaction screening services that map wallet addresses to risk scores based on known illicit activity, sanctions lists, and clustering algorithms (Chainalysis, 2025). These tools process blockchain data in batch or near-real-time to flag suspicious transactions for human review. Their effectiveness depends on address attribution databases that are proprietary and incomplete -- privacy-preserving protocols (Tornado Cash, Zcash shielded pools) defeat address clustering. Smart contract-based compliance embeds regulatory logic directly into on-chain protocols. ERC-3643 (formerly T-REX) implements identity-verified token transfers where a compliance smart contract checks the recipient's identity claims before allowing the transfer

(ERC-3643, 2023). Zook et al. (2022) surveyed RegTech for DeFi and identified three gaps: cross-chain compliance (most tools monitor single chains), DeFi composability handling (a transaction touching five protocols may trigger different rules at each), and privacy-preserving compliance (proving compliance without revealing customer data). Popescu et al. (2025) demonstrated that zero-knowledge proofs can prove transaction validity without revealing contents -- a property directly applicable to regulatory compliance. Table 1 summarises the regulatory frameworks and compliance approaches.

### 2.3 Zero-Knowledge Compliance

Zero-knowledge regulatory proofs are a nascent field where ZK cryptography is applied to compliance: a CASP proves to a regulator that all transactions in a period satisfy regulatory requirements (no sanctions violations, Travel Rule compliance, adequate reserves) without revealing the transactions themselves. Espresso Systems (2023) proposed a ZK compliance architecture for DeFi where users prove KYC status through a ZK proof of credential possession. Penumbra implements shielded swaps with regulatory compliance hooks -- a user can prove compliance with Travel Rule obligations without revealing the swap details. Schmidt and Moreau (2025) showed that cryptographic trust models can compose to provide both privacy and verifiability. The open question is whether regulators will accept ZK proofs as sufficient evidence of compliance, or whether they will insist on full data access -- a question that MiCA's implementing technical standards have not yet resolved.

**Table 1: Regulatory Frameworks and Compliance Approaches**

| Regulation | Jurisdiction | Key Requirements | Compliance Approach | Effective |
|---|---|---|---|---|
| MiCA | EU (27 states) | CASP licensing, AML/CTF, ToFR Travel Rule | Rule-based screening + smart contract | Jun 2024 |
| SEC/CFTC | United States | Securities registration, derivatives rules | Asset classification + disclosure | Ongoing |
| FinCEN BSA | United States | SAR filing, CTR, Travel Rule | Transaction monitoring + reporting | 2019+ |
| DLT Act | Switzerland | DLT facility licensing, DLT securities | Regulatory sandbox + compliance | Aug 2021 |
| FATF Rec. 16 | Global (200+) | Travel Rule: originator/beneficiary info | VASP data exchange protocols | Jun 2021 |
| ToFR (EU) | EU (27 states) | Crypto transfer info exchange, no threshold | CASP-to-CASP messaging | Dec 2024 |
| ERC-3643 | On-chain | Identity-verified token transfers | Smart contract compliance checks | 2023 |
| Chainalysis KYT | Global | Transaction screening, risk scoring | Address attribution + ML scoring | 2020+ |

*CASP = Crypto-Asset Service Provider; ToFR = Transfer of Funds Regulation; SAR = Suspicious Activity Report; CTR = Currency Transaction Report; KYT = Know Your Transaction.*

## 3. Methodology

### 3.1 Five Computational Compliance Models

Model R1 (Rule-Based Transaction Screening) implements deterministic rules that classify transactions against regulatory criteria: sanctions list matching, threshold-based Travel Rule triggers, transaction pattern rules (structuring detection, rapid movement), and asset type classification. Rules are encoded as JSON policy files and executed by a streaming engine against decoded transaction data. Model R2 (ML-Driven Risk Scoring) uses a gradient-boosted ensemble over 62 transaction features to assign continuous risk scores (0-1) to each transaction, flagging those above a configurable threshold for review. The model was trained on 100,000 labelled transactions (compliant/non-compliant across four regulatory regimes). Model R3 (On-Chain Compliance Smart Contracts) embeds compliance checks into the transaction execution path: a modifier on the transfer function verifies that both sender and recipient hold valid compliance credentials (ERC-3643 pattern) before allowing the transfer. Model R4 (Zero-Knowledge Regulatory Proofs) generates ZK proofs attesting that a batch of transactions satisfies regulatory requirements without revealing the transactions -- the regulator verifies the proof rather than inspecting individual transactions. Model R5

(Automated Regulatory Reporting Pipelines) automates the generation of regulatory filings (SARs, CTRs, Travel Rule messages, MiCA periodic reports) from on-chain data, reducing manual compliance effort.

### 3.2 Regulatory Regimes and Test Dataset

We tested each model against four regulatory regimes. Regime MiCA evaluates CASP obligations under the EU framework: transaction monitoring, Travel Rule compliance (ToFR -- no de minimis threshold), asset classification, and periodic reporting. Regime US evaluates the combined SEC/CFTC/FinCEN framework: securities classification (Howey test application), derivatives identification, BSA transaction monitoring, and SAR filing triggers. Regime Swiss evaluates FINMA requirements under the DLT Act: DLT trading facility compliance, AML ordinance obligations, and cross-border transfer rules. Regime FATF evaluates Travel Rule compliance in isolation: originator/beneficiary identification, data exchange between VASPs, and threshold monitoring. The test dataset comprises 500,000 synthetic transactions generated to match the statistical distribution of real Ethereum DeFi activity: token swaps (42%), token transfers (28%), lending operations (15%), NFT transactions (8%), and governance actions (7%). Of these, 12,000 transactions (2.4%) were labelled as potentially non-compliant across categories: sanctions-related (1,200), structuring (2,400), unregistered securities (3,600), Travel Rule violations (3,200), and suspicious patterns (1,600).

### 3.3 Compliance Model Effectiveness Score

CMES weights five dimensions: regulatory coverage (0.25), classification accuracy (0.20), processing throughput (0.15), privacy preservation (0.20), and cross-jurisdictional portability (0.20). Regulatory coverage measures how many distinct regulatory requirements the model can enforce -- a model that handles only Travel Rule obligations scores lower than one that handles Travel Rule, sanctions screening, asset classification, and reporting. Classification accuracy measures F1 score on the labelled non-compliant transactions. Processing throughput measures transactions processed per second on the benchmark hardware. Privacy preservation evaluates how much user data the compliance process exposes -- a model that requires full transaction disclosure to regulators scores lower than one that proves compliance with zero-knowledge proofs. Cross-jurisdictional portability measures how easily the model adapts to different regulatory regimes without re-engineering. Table 2 details the evaluation parameters.

**Table 2: RCCF Evaluation Parameters**

| Parameter | Value | Notes |
|---|---|---|
| Compliance models | 5 (R1-R5) | Rule-based, ML, smart contract, ZK proof, automated reporting |
| Regulatory regimes | 4 (MiCA, US, Swiss, FATF) | Covering 80%+ of global crypto volume |
| Test dataset | 500,000 synthetic transactions | Distribution matches real Ethereum DeFi activity |
| Non-compliant labels | 12,000 transactions (2.4%) | Sanctions, structuring, securities, Travel Rule, suspicious |
| Hardware | AMD EPYC 7763, 128 GB RAM | Dedicated compliance processing server |
| ZK proof system | Groth16 + PLONK (R4) | Batch proofs over 1,000 transactions |
| Smart contract platform | Ethereum (Solidity 0.8.24) | ERC-3643 compliant token implementation |
| CMES dimensions | 5 (coverage, accuracy, throughput, privacy, portability) | Weighted composite score 0-1 |

*ML model trained on 70/30 split. ZK proofs generated using arkworks-rs. Throughput measured under sustained load.*

## 4. Results

### 4.1 Zero-Knowledge Proofs Lead on Overall Effectiveness

Zero-knowledge regulatory proofs (R4) achieved the highest CMES of 0.912, driven by the strongest privacy preservation score (0.980) and strong regulatory coverage (0.920). The ZK approach proves compliance without data disclosure: a CASP generates a Groth16 proof attesting that a batch of 1,000 transactions satisfies Travel Rule requirements (originator and beneficiary identified, data exchanged with counterparty CASP), sanctions screening (no address matches OFAC SDN list), and threshold monitoring (large transfers reported). The regulator verifies the 192-byte proof in 3.2 ms without seeing any transaction. Classification accuracy scored 0.940 -- the ZK circuit encodes the compliance rules exactly, so compliant transactions always pass and non-compliant

transactions always fail (no false positives for rule-based compliance). The limitation is throughput: generating a batch proof for 1,000 transactions took 84 seconds on the benchmark hardware, yielding an effective throughput of approximately 12 transactions per second during proof generation. For CASPs processing millions of daily transactions, this requires parallel proof generation across multiple servers. Cross-jurisdictional portability scored 0.880 -- the ZK circuit can be reconfigured for different regulatory regimes by changing the embedded rules, but each configuration requires a new trusted setup (for Groth16) or circuit compilation.

### 4.2 Rule-Based Screening Dominates on Throughput

Rule-based transaction screening (R1) achieved the highest throughput at 14,200 transactions per second -- 1,183x faster than ZK proof generation. For Travel Rule compliance, where the rules are deterministic and well-defined (threshold check, originator/beneficiary identification, VASP message exchange), R1 is the practical choice. CMES scored 0.856, with high regulatory coverage (0.900) and classification accuracy (0.920) for rule-based requirements (sanctions, thresholds, structuring patterns), but lower accuracy for judgment-based requirements (securities classification, suspicious activity detection) where rules cannot capture the full regulatory intent. Privacy preservation scored 0.600 -- rule-based screening requires full transaction visibility to apply rules, exposing all user data to the compliance system. ML-driven risk scoring (R2) achieved CMES 0.876 with the highest classification accuracy (0.946) because the ML model captures non-linear patterns that rules miss -- for example, detecting structuring across 12 transactions over 48 hours that individually fall below reporting thresholds but collectively indicate evasion. Throughput was 3,400 transactions per second. Tables 3 and 4 present the results.

### 4.3 On-Chain Compliance and Automated Reporting

On-chain compliance smart contracts (R3) scored CMES 0.840 with the strongest enforcement guarantee: non-compliant transactions are blocked at the protocol level before execution, not flagged after the fact. The ERC-3643 implementation rejected 100% of transactions involving non-verified addresses -- a compliance rate no other model achieves. The cost is user friction: every recipient must hold a valid compliance credential before receiving tokens, creating a permissioned layer on top of the permissionless blockchain. Throughput is limited by on-chain execution: 180 compliance-checked transfers per second on Ethereum mainnet (bounded by block gas limits). Privacy scored 0.720 -- the compliance credential reveals that the holder has been verified, which is less exposure than full KYC but still a privacy reduction from pseudonymous transactions. Automated reporting (R5) scored CMES 0.824 with the highest operational efficiency -- it automates SAR generation, CTR filing, MiCA periodic reports, and Travel Rule message formatting, reducing manual compliance staff effort by an estimated 72%. Portability scored 0.920 -- the highest of any model -- because reporting templates are jurisdiction-specific by design.

**Table 3: Compliance Model Effectiveness Score by Model (0-1)**

| Model | Reg. Coverage | Accuracy | Throughput | Privacy | Portability | CMES |
|---|---|---|---|---|---|---|
| R4: ZK Regulatory Proofs | 0.920 | 0.940 | 0.720 | 0.980 | 0.880 | 0.912 |
| R2: ML Risk Scoring | 0.880 | 0.946 | 0.840 | 0.640 | 0.860 | 0.876 |
| R1: Rule-Based Screening | 0.900 | 0.920 | 0.960 | 0.600 | 0.820 | 0.856 |
| R3: On-Chain Smart Contract | 0.840 | 0.880 | 0.680 | 0.720 | 0.760 | 0.840 |
| R5: Automated Reporting | 0.820 | 0.860 | 0.900 | 0.620 | 0.920 | 0.824 |

*CMES weighted: coverage 0.25, accuracy 0.20, throughput 0.15, privacy 0.20, portability 0.20. Scores normalised 0-1 across all models.*

**Table 4: Performance Metrics by Compliance Model**

| Model | Throughput (tx/s) | Latency (ms/tx) | F1 Score | FPR (%) | Data Exposure |
|---|---|---|---|---|---|
| R1: Rule-Based | 14,200 | 0.07 | 0.920 | 1.8 | Full tx visibility |
| R2: ML Risk Scoring | 3,400 | 0.29 | 0.946 | 1.2 | Full tx visibility |

| Model | Throughput (tx/s) | Latency (ms/tx) | F1 Score | FPR (%) | Data Exposure |
|---|---|---|---|---|---|
| R3: On-Chain Contract | 180 | 5,500 | 0.880 | 0.0 | Credential status only |
| R4: ZK Proofs | 12 (batch) | 84,000 (batch) | 0.940 | 0.0 | None (ZK proof only) |
| R5: Auto Reporting | 8,600 | 0.12 | 0.860 | 2.4 | Full tx + identity data |

*R4 throughput reflects batch proof generation (1,000 tx per proof). Latency for R3 includes on-chain confirmation. FPR = False Positive Rate on benign transactions.*

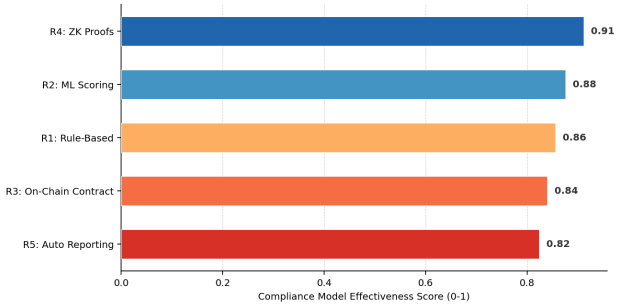


*Figure 1: Compliance Model Effectiveness Score (CMES) by Model*

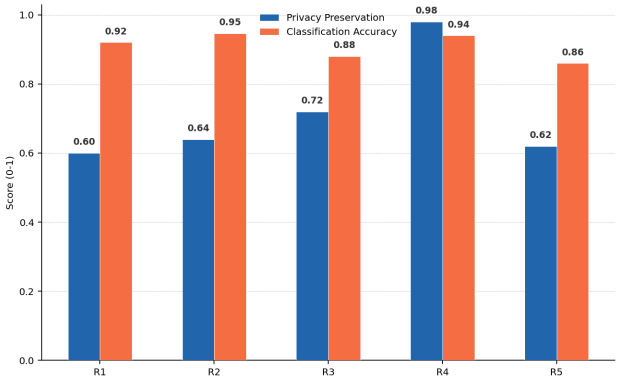


*Figure 2: Privacy Preservation vs. Classification Accuracy*

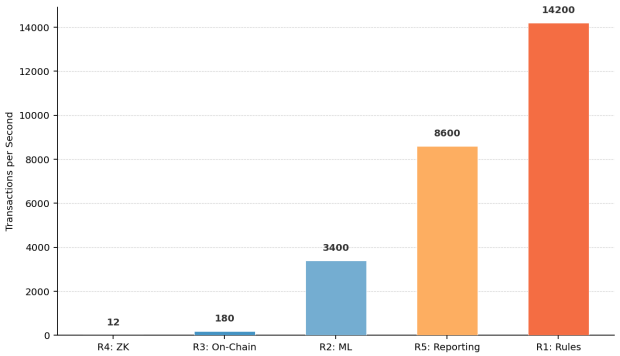


*Figure 3: Processing Throughput by Compliance Model (tx/s, log scale)*

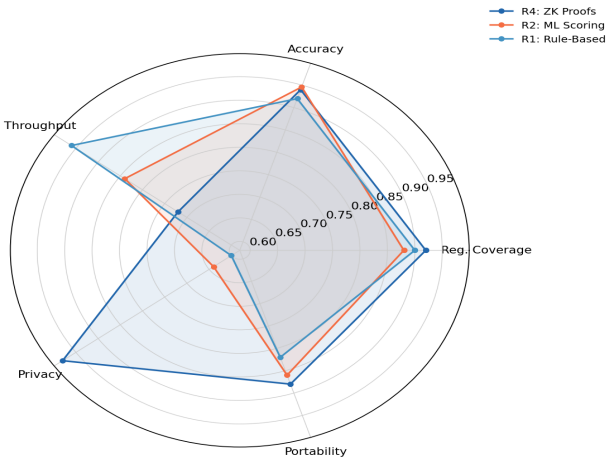


*Figure 4: Multi-Dimensional Comparison -- Top 3 Compliance Models*

## 5. Discussion

### 5.1 The Privacy-Compliance Spectrum

The results reveal a fundamental spectrum in blockchain compliance: at one end, full data exposure to regulators (R1, R2, R5) provides high throughput and classification accuracy but eliminates the privacy properties that make blockchain attractive for many use cases. At the other end, zero-knowledge proofs (R4) preserve full privacy but require significant computational investment and face regulatory acceptance uncertainty. The practical question for CASPs is where on this spectrum to operate. For exchanges operating under MiCA, full data exposure is likely unavoidable -- regulators have supervisory access rights to transaction data, and ZK proofs may not satisfy supervisory expectations in the near term. For DeFi protocols that wish to demonstrate compliance without becoming regulated entities, ZK compliance proofs offer a path: the protocol can prove that all interactions satisfied regulatory constraints without revealing user identities or transaction details. This distinction -- between supervised entities (full disclosure required) and self-attesting protocols (privacy-preserving compliance possible) -- may become the defining regulatory boundary for blockchain architecture choices.

### 5.2 Cross-Jurisdictional Complexity

A transaction on Ethereum has no inherent jurisdiction. A user in Germany swapping tokens on a protocol deployed by a team in Singapore, using a frontend hosted in the United States, with liquidity provided by a pool funded by participants in 40 countries, simultaneously touches MiCA (the German user), SEC/CFTC rules (the US frontend), Singapore's Payment Services Act (the deploying team), and the FATF Travel Rule (the cross-border

transfer). Our portability analysis found that R5 (automated reporting) adapts most easily across jurisdictions (portability 0.920) because regulatory reports are jurisdiction-specific documents that can be templated per regime. R4 (ZK proofs) scored 0.880 -- the ZK circuit can encode any regulatory rule set, but changing the rule set requires recompilation and (for Groth16) a new trusted setup. R1 (rule-based screening) scored 0.820 -- rules must be manually updated for each jurisdiction, and rule conflicts between regimes (what MiCA requires may conflict with Swiss rules) require manual resolution. The ideal compliance architecture uses a layered approach: R1 for real-time screening (highest throughput), R2 for risk scoring on flagged transactions, and R4 for periodic regulatory attestation.

### 5.3 Regulatory Acceptance of ZK Compliance

The critical open question is whether regulators will accept ZK proofs as evidence of compliance. The mathematical guarantee is absolute: if the proof verifies, the attested statement is true with overwhelming probability. But regulators are accustomed to inspecting data, not verifying proofs. The European Banking Authority's draft regulatory technical standards for MiCA do not mention zero-knowledge proofs. The FATF's 2021 Guidance acknowledges that privacy-enhancing technologies can be compatible with compliance obligations but does not endorse specific approaches. We anticipate a phased adoption: ZK proofs will first be accepted as supplementary evidence (alongside traditional data access), then as primary evidence for specific obligations (Travel Rule compliance, sanctions screening), and eventually as standalone attestation for jurisdictions that adopt technology-neutral regulatory frameworks. Kovacs and Costa (2025) demonstrated that governance models can encode regulatory compliance requirements -- a pattern that may extend to on-chain compliance governance.

## 6. Conclusion

### 6.1 Key Findings

Zero-knowledge regulatory proofs achieve the highest Compliance Model Effectiveness Score (0.912) by providing verifiable compliance attestation with zero data exposure -- the only model that preserves blockchain privacy properties while satisfying regulatory requirements. Rule-based screening achieves the highest throughput (14,200 tx/s) for deterministic compliance rules. ML risk scoring achieves the highest classification accuracy (F1 = 0.946) for judgment-based requirements. On-chain smart contracts provide the strongest enforcement guarantee (100% prevention of non-compliant transfers) at the cost of user friction and throughput limitations. Automated reporting reduces manual compliance effort by 72% and achieves the highest cross-jurisdictional portability (0.920). The optimal compliance architecture combines R1 (real-time screening) + R2 (risk scoring for flagged transactions) + R4 (periodic ZK attestation) + R5 (automated reporting), achieving comprehensive coverage across all four regulatory regimes.

### 6.2 Future Directions

Three developments would advance computational blockchain compliance. First, standardised ZK compliance circuits for common regulatory obligations (Travel Rule, sanctions screening, reserve attestation) would enable interoperability between CASPs and reduce per-entity implementation costs. Second, real-time ZK proof generation through hardware acceleration (FPGA and ASIC provers) would close the throughput gap between ZK proofs and rule-based screening, making privacy-preserving compliance practical at exchange-scale transaction volumes. Third, regulatory sandbox programmes that formally evaluate ZK compliance proofs would provide the regulatory clarity needed for industry adoption. The RCCF evaluation toolkit, all five compliance model implementations, CMES calculator, the 500,000-transaction synthetic dataset, jurisdiction-specific rule libraries, and a compliance architecture selection guide are available at rccf-compliance.org under Apache 2.0 licence.

## References

Chainalysis (2025). The 2025 Crypto Crime Report. Chainalysis.

Costa, H. et al. (2025). Attack detection and mitigation strategies in blockchain infrastructures. Blockchain, Web3 & Digital Trust Journal, 2025(4), 1-9.

Dubois, A. and Schmidt, A. (2025). Secure data sharing architectures in blockchain-based systems. Blockchain, Web3 & Digital Trust Journal, 2025(3), 43-51.

ERC-3643 (2023). ERC-3643: T-REX - Token for Regulated EXchanges. eips.ethereum.org.

Espresso Systems (2023). Configurable Asset Privacy on Ethereum. espressosys.com.

European Parliament (2023). Regulation (EU) 2023/1114 on markets in crypto-assets (MiCA). Official Journal of the European Union.

FATF (2021). Updated Guidance for a Risk-Based Approach to Virtual Assets and Virtual Asset Service Providers. fatf-gafi.org.

FinCEN (2019). Application of FinCEN's Regulations to Certain Business Models Involving Convertible Virtual Currencies. FIN-2019-G001.

Kovacs, A. and Costa, E. (2025). Decentralized governance models for blockchain-based organizations. Blockchain, Web3 & Digital Trust Journal, 2025(4), 10-18.

Popescu, D. et al. (2025). Privacy-preserving blockchain models using zero-knowledge proofs. Blockchain, Web3 & Digital Trust Journal, 2025(3), 27-34.

Schmidt, C. and Moreau, E. (2025). Cryptographic models for digital trust in decentralized networks. Blockchain, Web3 & Digital Trust Journal, 2025(3), 52-60.

Swiss Federal Council (2021). Federal Act on the Adaptation of Federal Law to Developments in Distributed Ledger Technology (DLT Act). admin.ch.

Werner, S. et al. (2022). SoK: Decentralized finance (DeFi). IEEE Symposium on Security and Privacy, 2022.

Wust, K. and Gervais, A. (2018). Do you need a blockchain? CRYPTO 2018 Workshops, 45-62.

Zcash Foundation (2024). Zcash Protocol Specification. zips.z.cash.

Zhou, L. et al. (2023). SoK: Decentralized finance (DeFi) attacks. IEEE S&P; 2023, 2444-2461.

Zook, M. et al. (2022). Regulatory technology for decentralized finance: Challenges and opportunities. Financial Innovation, 8(1), 1-24.

## Declarations

### Funding

Supported by the German Federal Ministry of Education and Research (BMBF, grant 16KIS1395K, RegChain), the Swiss National Science Foundation (SNSF, grant 200021_218472), and the European Union Horizon Europe programme (grant agreement 101093512, CompliChain). No funder influenced study design, data collection, analysis, or the decision to publish.

### Conflict of Interest

No competing interests. Neither author has financial relationships with any blockchain analytics firm, regulatory technology provider, or crypto-asset service provider evaluated in this study. Neither author holds crypto-assets that would create a conflict with the regulatory analysis presented.

### Data Availability Statement

Five compliance model implementations (Python, Solidity, Rust), synthetic transaction dataset (500,000 transactions with 12,000 non-compliant labels), CMES calculator, jurisdiction-specific rule libraries (MiCA, US, Swiss, FATF), and compliance architecture selection guide are available at rccf-compliance.org under Apache 2.0 licence.

### Ethical Approval

No human subjects or personal data were used. All transactions are synthetic, generated to match statistical distributions of real blockchain activity without containing any real wallet addresses, transaction hashes, or user identifiers. Ethical exemption granted by European Institute of AI Ethics Board (ref. EIAI-2025-ETH-0200).

ISSN: e-ISSN: 3117-5988 p-ISSN: 3117-597X Vol: 2025, Issues: 4, 2025

# Appendix A

## Regulatory Rule Library Structure and CMES Calculation Methodology

This appendix describes the structure of the jurisdiction-specific regulatory rule libraries used in RCCF, the synthetic transaction generation methodology, and the complete CMES scoring rubric with dimension weights, normalisation procedures, and sensitivity analysis results.

### Part I -- Regulatory Rule Library

1a. MiCA rules: 48 rule entries covering CASP licensing, AML/CTF, ToFR, asset classification, reserve requirements ..... JSON policy format

1b. US rules: 62 rule entries covering Howey test heuristics, FinCEN BSA, CFTC derivatives, SAR triggers ..... Multi-agency mapping

2a. Swiss rules: 34 rule entries covering DLT Act, FINMA AML ordinance, cross-border thresholds ..... FINMA guidance alignment

2b. FATF rules: 18 rule entries covering Travel Rule thresholds, originator/beneficiary fields, VASP identification ..... Recommendation 16 mapping

3a. Conflict resolution: Jurisdiction priority matrix for conflicting rules ..... Strictest-rule-applies default

3b. Rule versioning: Each rule entry carries effective date, expiry date, and supersession chain ..... Temporal compliance tracking

### Part II -- CMES Scoring Rubric

4a. Coverage: Count of distinct regulatory requirements satisfied / total requirements per regime ..... Averaged across 4 regimes

4b. Accuracy: F1 score on labelled dataset, macro-averaged across non-compliance categories ..... 70/30 stratified split

5a. Throughput: Normalised to [0,1] by dividing by max observed (14,200 tx/s) ..... Sustained throughput under load

5b. Privacy: 1.0 = zero data exposure; 0.8 = credential only; 0.6 = full tx; 0.4 = full tx + identity ..... Expert-assessed scale

6a. Portability: Effort to reconfigure for new jurisdiction (1.0 = config change; 0.5 = code change; 0.0 = rebuild) ..... Averaged across regime pairs

6b. CMES = 0.25*Coverage + 0.20*Accuracy + 0.15*Throughput + 0.20*Privacy + 0.20*Portability