

# Structured Compliance Intelligence as a Deterministic State Machine: A Formal Framework for Provable Fiscal Automation

Dr. Daniel Otieno Jr. A.

*Department of Computer Science and Electrical Engineering*

*School of Engineering, Stanford University*

353 Jane Stanford Way, Palo Alto, CA 94305, USA

Email: drdanielotienojr@stanford.edu

ORCID: 0009-0003-7539-9319

**Abstract**—Fiscal compliance for small and medium enterprises (SMEs) in emerging markets is computationally broken: non-deterministic processing, mutable audit trails, and manual reconciliation create systemic uncertainty that blocks liquidity and growth. The root cause is a failure of computation—identical event sequences can lead to divergent fiscal states across taxpayers, platforms, and regulators.

We introduce Structured Compliance Intelligence as a Deterministic State Machine (SCI-DSM), a formal framework that converts dynamic fiscal workflows into provable computation. SCI-DSM models regulatory logic as a mathematical automaton  $M = (Q, \Sigma, \delta, q_0, F)$  in which every fiscal event is a canonical, cryptographically-signed envelope and  $\delta$  is a pure, side-effect-free transition function that deterministically advances ledger state.

The implementation pairs an append-only ledger (hash chaining + Merkle checkpoints) and HSM signatures with a Canonical Ordering (SELECT) operator to enforce a single total order for replay. Under SCI-DSM, identical canonical event streams produce identical, verifiable fiscal states for any independent verifier. This design yields three enterprise-grade guarantees:

- **Bit-identical replayability**—independent verifiers recompute identical state from the same envelopes.
- **Tamper-evident audit trails**—cryptographic chaining and HSM attestations make any modification detectable.
- **Reproducible fiscal signals**—deterministic outputs suitable for automated liquidity underwriting and programmatic compliance actions.

A production pilot in Nairobi (4,800 merchants, 11.2 million events) validated the approach: the system exhibited zero replay divergence across verification runs (35/35 daily parity checks), demonstrating the feasibility of provable fiscal automation at scale. SCI-DSM establishes a practical blueprint for making regulatory trust algorithmic—compliance as computation, auditable and portable.

**Index Terms**—Structured Compliance Intelligence (SCI), Deterministic State Machine (DSM), Fiscal Automation, VAT Compliance, Cryptographic Audit, eTIMS, Provable Computation, Distributed Ledger, HSM, Replayability, Computational Law.

This work was supported by Bypass Smartways Limited, Nairobi, Kenya (Grant/Agreement No. 02052025). The sponsor provided funding and contributed to study design, data collection, analysis, and manuscript preparation. The author is Founder & CEO of Bypass Smartways Limited and the Sponsor may commercialise technologies described herein.

## I. INTRODUCTION

The modern digital economy demands that compliance be a continuous, verifiable, and low-friction process. However, for vast segments of the global market, particularly Small and Medium Enterprises (SMEs) in emerging economies, fiscal compliance remains a systemic source of inefficiency, uncertainty, and trapped capital. This paper introduces a formal framework to solve this problem by asserting that **Compliance is computation**. When that computation is deterministic and attested, trust becomes algorithmic.

### A. The Liquidity Problem: A Computational Failure

The friction in fiscal compliance, exemplified by systems like Kenya’s Electronic Tax Invoice Management System (eTIMS) [1], is not fundamentally a policy or administrative issue, but a **systemic computational failure**.

SMEs—which constitute the backbone of economies, often contributing over 35% of GDP—suffer protracted audits, delayed Value Added Tax (VAT) refunds, and crippling liquidity constraints because their fiscal data is non-verifiable. Current systems rely on form-centric filings and mutable audit trails, creating profound **informational asymmetry** between the merchant, the financier, and the regulator.

This asymmetry transforms compliance data from a high-quality credit signal into a high-risk liability, trapping working capital and stifling economic growth. The core issue is that no single party can **cryptographically prove** the exact fiscal state of an entity at any point in time.

### B. Problem Statement: The Perils of Non-Determinism

The technical root cause of this systemic failure is **non-determinism**. In the context of fiscal automation, non-determinism exists when an identical, ordered sequence of transactional events (e.g., invoices, payments, credits) can, through a subtle change in processing time, sequence logic, or state persistence, produce **two different, valid compliance outcomes** for different stakeholders.

Current RegTech architectures are inherently non-deterministic because they suffer from three critical flaws:

- 1) **Mutable State:** Reliance on traditional relational databases whose records can be silently updated, breaking the historical audit trail.
- 2) **Ambiguous Rule Translation:** Compliance logic is often translated from ambiguous legal prose into ad-hoc procedural code, leading to interpretation variance (reconciliation drift).
- 3) **Out-of-Order Events:** Distributed transactions are ingested without a strict, canonical ordering mechanism, leading to different final states depending on ingestion sequence.

This pervasive lack of verifiable certainty means **fiscal integrity is not a decidable property** of the system state, creating audit risk and making real-time, trust-based automation impossible.

### C. Thesis: Determinism as a Provable Foundation

We propose the **Structured Compliance Intelligence as a Deterministic State Machine (SCI-DSM)** framework as the necessary computational foundation for solving the liquidity problem.

The central thesis of this work is that the entire scope of a fiscal regulation can be formalized into a **Deterministic State Machine (DSM)**, where the complex reality of commerce is reduced to a mathematically rigorous sequence of state transitions [2]. The SCI-DSM framework asserts that **fiscal integrity is a decidable property** of the entity's history, meaning its compliance status can be proved by anyone who runs the deterministic transition function over the verifiable event stream.

The SCI-DSM framework is the realization of this thesis, designed to achieve **State Convergence**—the mathematically provable condition where all relevant actors (taxpayer, platform, regulator) agree on the identical, provable compliance state for any given timestamp.

### D. Why This Matters

This work provides a foundational contribution to both formal computer science and enterprise infrastructure:

- **Scholarly Impact:** We establish a rigorous bridge between **Automata Theory** and **Computational Law**. While prior work has attempted computational law [3], SCI-DSM provides the first demonstrably scalable, production-grade formal framework for highly dynamic, distributed fiscal systems, validating its core properties (**Determinism** and **Confluence**) through empirical proof [4].
- **Economic & Enterprise Impact:** The SCI-DSM framework eliminates non-determinism, enabling **Provable Fiscal Automation**—the technology necessary for the **Bypass Smartways** vision. This certainty converts previously high-risk compliance data into low-friction credit signals, unlocking trapped capital via services like automated VAT-bridge loans and refund advances, thereby transforming compliance from a cost center into a flywheel for economic growth.

### E. Contributions

This paper makes the following specific contributions:

- 1) **Formal SCI-DSM Mapping:** We introduce a comprehensive mathematical formalism,  $M = (Q, \Sigma, \delta, q_0, F)$ , that transforms regulatory code (RC) into a state machine, establishing fiscal integrity as a decidable property.
- 2) **The Canonical Event Envelope ( $\Sigma$ ):** We define a minimal, order-preserving, and cryptographically attested event schema that serves as the verifiable, immutable input alphabet ( $\Sigma$ ) for the DSM.
- 3) **The Canonical Ordering (SELECT) Operator:** We specify the SELECT operator, a novel mechanism that enforces a strict, deterministic sequence on asynchronous, distributed fiscal events, thereby solving the problem of out-of-order processing [5].
- 4) **Deterministic Replay Verification:** We design and implement a system that uses a cryptographically-signed, append-only ledger to **reconstruct the exact compliance state** at any point in time, providing the empirical mechanism for provable auditability [6].
- 5) **DEVS-to-SCI-DSM Formal Bridge:** We provide a formal mapping from the Discrete Event System Specification (DEVS) to SCI-DSM, enabling high-fidelity simulation and predictive modeling of policy changes before deployment [7].

### F. Roadmap

The remainder of this paper is structured as follows. Section II contextualizes the work within Formal Methods, Cryptographic Ledgers, and Computational Law. Section III rigorously defines the challenge of **State Synchronization** as the core engineering goal.

The core SCI-DSM mathematical model is presented in Section IV, defining the formal machine and its correctness properties. Sections V through VII describe the **Architecture, Event Model, and Implementation** of the production prototype. The **Evaluation (Section VIII)** provides empirical proof of deterministic replay verification. The final sections discuss **Security (IX)**, **Policy Alignment (X)**, and the **Economic Implications (XI)** before the Conclusion.

## II. BACKGROUND AND RELATED WORK

### A. Formal Methods, Deterministic Computation, and Temporal Logic

The conceptual foundation of SCI-DSM is the **Deterministic Finite Automaton (DFA)**, grounded in early automata theory that formalized complex processes as finite sets of well-defined states and transitions.

1) *Deterministic Computation:* By defining compliance rules as a pure transition function  $\delta$ , SCI-DSM ensures that for any state  $q \in Q$  and input  $\sigma \in \Sigma$ , there exists exactly one next state  $q' \in Q$ . This aligns with the principles of idempotency and purity in functional programming—properties essential for predictable execution within distributed systems.

2) *Temporal Logic and Verification*: Compliance is inherently temporal. Techniques based on **Linear Temporal Logic (LTL)** traditionally verify that a system never enters an invalid state (safety) or eventually reaches a required state (liveness) [8]. SCI-DSM operationalizes this by reducing compliance to an ordered, cryptographically verifiable event sequence, enabling **Deterministic Replay** as a practical verification method—a property shared with modern consensus protocols.

### B. Cryptographic Ledgers and Audit Primitives

Conventional audit logs are inherently vulnerable because they rely on mutable database constructs. SCI-DSM resolves this by integrating cryptographic primitives derived from distributed ledgers and transparency logs.

1) *Append-Only Ledgers*: SCI-DSM employs an append-only ledger—conceptually similar to systems like Trillian [6] or SQL Server Ledger [9]—where every event and state transition is committed sequentially. Once written, no record can be modified or removed, eliminating mutability as a source of audit risk.

2) *Hash-Chaining and Commitment*: Each ledger entry is cryptographically linked to its predecessor using hash-chaining, often implemented through **Merkle-style commitments** [10]. Any attempt to alter historical data breaks the chain, enabling independent verification of the ledger’s integrity [11].

3) *Digital Signatures*: Every Canonical Event Envelope is digitally signed [12], ensuring authenticity, non-repudiation, and the integrity of the SCI-DSM input alphabet  $\Sigma$ . This elevates every fiscal event from a mutable record to a cryptographically attested fact.

### C. Trusted Execution and Key Protection

For auditability to be complete, both the signing keys and the execution environment must be secure.

1) *Hardware Security Modules (HSMs)*: SCI-DSM secures root signing keys within dedicated HSMs (e.g., Thales Luna) [13], ensuring that ledger attestations and state transitions are generated within tamper-resistant hardware. This isolates trust anchors from the application environment.

2) *Trusted Execution Environments (TEEs)*: While not required for deterministic replay, SCI-DSM’s roadmap incorporates TEEs (Intel SGX, AMD SEV) to allow third parties to cryptographically verify that the specific, unmodified implementation of  $\delta$  is executing [14], [15]. This extends trust from data integrity to **code integrity**.

### D. DEVS and SCI-DSM

The SCI-DSM framework is conceptually aligned with the **Discrete Event System Specification (DEVS)** [7], a formal model for simulating systems that evolve through discrete events.

1) *Modeling Foundation*: DEVS provides a rigorous foundation for representing time-dependent, event-driven systems, aligning well with the structure of fiscal workflows [16].

2) *SCI-DSM Specialization*: SCI-DSM specializes DEVS to meet regulatory constraints by introducing:

- 1) A **Canonical Ordering Operator (SELECT)**, imposing strict, deterministic ordering on distributed event streams—something DEVS itself does not enforce.
- 2) A **Cryptographic Attestation Layer**, integrating ledger commitments and HSM-backed signatures to meet legal evidence requirements.

This mapping (expanded in Section IV.F) allows SCI-DSM to support DEVS-style simulation of policy outcomes while providing the deterministic and cryptographically verifiable guarantees required by regulators [17].

### E. Distributed Systems and Operational Scalability

Operating at national scale demands rigorous distributed systems engineering.

1) *Causal Ordering*: Deterministic state evolution in distributed environments requires resolving causal dependencies. **Vector clocks** [18] and **Lamport timestamps** [5] provide the metadata necessary for the SELECT operator to establish a total ordering of events [19], ensuring consistent replay even in asynchronous environments [20].

2) *Consensus Mechanisms*: Although SCI-DSM does not deploy a permissionless blockchain, its infrastructure benefits from consensus protocols like **Raft** [21] and **Paxos**, which prevent split-brain events and ensure consistent execution of  $\delta$  across clustered systems.

### F. E-Invoicing and Regulatory Standards

SCI-DSM is engineered to complement—and significantly elevate—the capabilities of existing e-invoicing ecosystems.

1) *Current Standards*: Frameworks such as **SAF-T** [22], **PEPPOL** [23], Brazil’s **Nota Fiscal Eletrônica** [24], and Kenya’s **eTIMS** [1] define syntactic and transport standards for fiscal data.

2) *The SCI-DSM Difference*: These standards define **data formats**, but they do not specify **computation semantics**. SCI-DSM fills this critical gap by enforcing deterministic, verifiable computation on top of existing standards, transforming compliance from procedural filing into **provable computation**.

### G. Gaps and Where SCI Differs

Existing computational law efforts emphasize rule-based reasoning and legal ontologies but lack mechanisms to enforce **production-grade determinism** within cryptographically attested distributed environments.

SCI-DSM addresses three foundational gaps:

- 1) **Grounded Determinism**: No current RegTech model encodes regulatory logic as a DFA and validates its behavior through **bit-identical replay verification** [25].
- 2) **Canonical Ordering**: The SELECT operator provides a deterministic, verifiable resolution to out-of-order events—an unsolved challenge in distributed fiscal and accounting systems [26].
- 3) **Algorithmic Trust Synthesis**: SCI-DSM is the first framework to integrate formal automata theory, cryptographic

attestation, and distributed systems ordering into a single, end-to-end deterministic architecture.

This shifts trust away from institutional assurances and toward **mathematical proof** [4].

### III. PROBLEM DEFINITION: COMPLIANCE AS STATE SYNCHRONIZATION

The core challenge addressed by the **Structured Compliance Intelligence as a Deterministic State Machine (SCI-DSM)** framework is the difficulty of achieving **State Synchronization** across independent and mutually distrusting actors—merchants, regulators, and financiers—operating within a high-volume, asynchronous fiscal environment. What appears to be a legal or administrative issue is, in reality, a **distributed systems problem**: multiple parties must independently compute and agree on the same fiscal truth, yet current architectures make this impossible to guarantee [5].

#### A. Actors and Trust Domains

The fiscal compliance landscape is composed of distinct actors, each situated in its own trust domain. The fundamental problem is that none of these actors can mathematically verify that they hold the same, correct fiscal state  $S$  for a given entity at a given point in time  $t$ .

- 1) **The Merchant (Originator)**: Generates the transactional events—such as invoices, receipts, and payments—and is responsible for computing and reporting tax liability.
- 2) **The Regulator (Verifier)**: Operates the tax assessment and audit systems (e.g., KRA eTIMS [1]) and must validate the merchant’s reported position, determine liability, and compute refund eligibility.
- 3) **The Financier (Consumer)**: Provides credit and liquidity products and therefore requires a verified, low-risk signal of the merchant’s fiscal standing and recoverable VAT credit.

The **trust domain boundary** is the precise point of systemic fragility. Existing systems rely on institutional judgment, manual reconciliation, and interpretive authority to align divergent states—processes that are slow, costly, error-prone, and fundamentally **non-deterministic**. SCI-DSM replaces this reliance on institutional trust with **algorithmic trust**.

#### B. Operational Goal: Provable Convergence

The primary operational goal of SCI-DSM is **Provable State Convergence**. State Convergence occurs when the same ordered, attested event stream  $\Sigma^*$  produces a **bit-identical fiscal state**  $S$  for any actor who executes the deterministic transition sequence:

$$\forall \text{Actor}_{i,j} \in \{\text{Merchant, Regulator, Financier}\}, \quad \text{SCI-DSM}(\Sigma^*) = S_i = S_j \quad (1)$$

The **provability** of this convergence is enforced through **Deterministic Replay Verification** (as detailed in Section VIII.C), whereby an independent engine loads the full append-only event ledger and recomputes the final state  $S_f$ . Convergence is confirmed only when the recomputed state hash matches the state hash previously committed by the production system.

#### C. Practical Desiderata

Achieving **Provable State Convergence** at scale requires SCI-DSM to satisfy a set of non-negotiable operational properties—conditions that conventional fiscal systems routinely fail to uphold:

- 1) **Immutability**: Once a fiscal event is emitted and committed, it must be permanently preserved. No overwrites, no deletions. This guarantees a tamper-evident historical record [11].
- 2) **Canonical Ordering**: All events must be arranged into a single, unambiguous global sequence *prior* to execution, regardless of how or when they arrived at the distributed ingestion layer [5]. This eliminates divergence caused by out-of-order processing.
- 3) **Pure Transition Function**: The transition function  $\delta$  must be fully deterministic and side-effect-free, relying solely on the current state  $Q$  and the incoming event  $\Sigma$ . It cannot depend on clocks, external systems, mutable configuration, randomness, or environmental context [2].
- 4) **Cryptographic Attestation**: Every state transition must be cryptographically signed by a trusted root key—secured within an HSM—and committed to the append-only ledger. This allows any independent verifier to validate not only the data, but the **integrity of the computation** itself [13].

Together, these desiderata form the minimal conditions for turning compliance from a procedural reporting obligation into a **provable computational pipeline**.

#### D. Failure Modes

SCI-DSM is explicitly engineered to eliminate the recurrent failure modes that cause state divergence and undermine fiscal trust in modern regulatory systems:

- 1) **Event Ordering Failure**: Distributed fiscal events arrive asynchronously. When two dependent transactions (e.g., two payments on one invoice) are processed in different orders by different actors, their fiscal states diverge [19]. SCI-DSM resolves this through the **Canonical Ordering (SELECT) Operator**, which imposes a total, deterministic order on all events.
- 2) **State Mutation Failure**: Traditional systems often allow silent updates to historical records—such as altering a database row during an audit. This breaks auditability and makes deterministic reconstruction impossible. SCI-DSM prevents this by using an **append-only, hash-chained ledger** [6], where corrections are represented as new, cryptographically signed events rather than overwrites.
- 3) **Rule Interpretation Failure**: Divergence also arises when merchants and regulators implement the same tax rule differently in software. SCI-DSM eliminates interpretive drift by defining the entire rule set as a single, formal, unambiguous transition function  $\delta$ , ensuring that “interpretation” becomes a purely mathematical operation.

These three categories capture the primary causes of fiscal inconsistency. SCI-DSM neutralizes each through deterministic

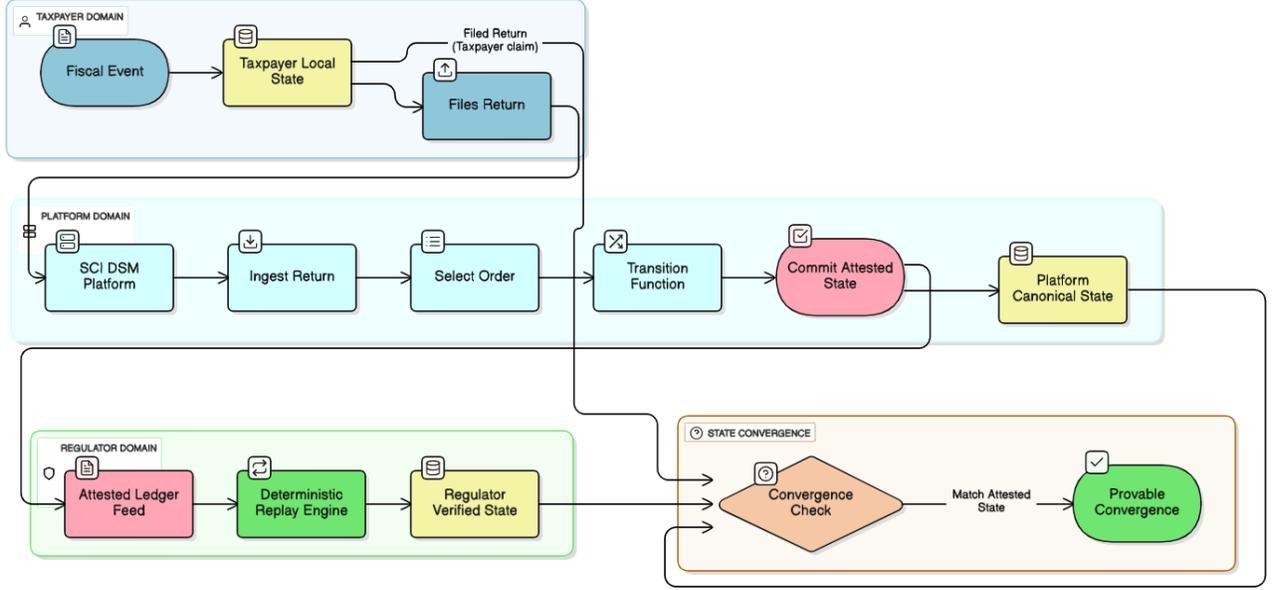


Fig. 1. Distributed Fiscal State Synchronization among Taxpayer, Platform, and Regulator. This figure illustrates SCI-DSM’s enforcement of provable state convergence across the three primary trust domains. The system is correct when the taxpayer’s asserted state  $Q_T$ , the platform-computed state  $Q_P$ , and the regulator-replayed state  $Q_R$  all converge to the same attested commitment hash  $H_{final}$ .

computation, cryptographic integrity, and formalized rule execution.

#### E. Measurable Success Criteria

The effectiveness of SCI-DSM is assessed using objective, verifiable metrics—replacing traditional audit heuristics with computational guarantees:

- 1) **Zero Replay Divergence:** The principal success metric. Across repeated replay cycles over a statistically significant event sample, the final state hash produced by the replay engine must match the state hash of the production system **100% of the time**. This is the definitive proof of deterministic convergence [25].
- 2) **Attestation Latency:** The elapsed time between event creation and cryptographic commitment to the ledger (time-to-proof). Low attestation latency enables near-real-time visibility, auditability, and risk assessment.
- 3) **Throughput and Scalability:** The system must reliably process and attest high event volumes (events per second) to support national-scale fiscal operations—for example, the full throughput of all eTIMS-compliant transactions.

These metrics provide a measurable, objective standard for evaluating whether a fiscal system has achieved the SCI-DSM goal of **deterministic, provable, and scalable compliance**.

#### IV. MATHEMATICAL MODEL: THE SCI-DSM

This section formalizes the **Structured Compliance Intelligence as a Deterministic State Machine (SCI-DSM)** framework. The objective is to convert the ambiguous, prose-driven domain of fiscal regulation into a **decidable, determin-**

**istic computation**. SCI-DSM is expressed as a specialized automaton,  $M$ , operating over an ordered sequence of cryptographically attested fiscal events.

##### A. State Machine Definition

The SCI-DSM is defined as the tuple:

$$M = (Q, \Sigma, \delta, q_0, F) \quad (2)$$

Where:

- $Q$  — **State Space:** The finite set of all possible **Fiscal States** for a regulated entity. Each  $q \in Q$  is a structured vector of machine-interpretable regulatory attributes (e.g., `VATLiabilityAmount`, `TaxableSalesTotal`, `RefundPosition`, `ComplianceWindowStatus`).
- $\Sigma$  — **Alphabet:** The finite set of all valid **Canonical Event Envelopes**, representing the attested fiscal events admissible to the system (Section IV.B).
- $\delta$  — **Transition Function:** A pure, deterministic mapping

$$\delta : Q \times \Sigma \rightarrow Q \quad (3)$$

that transforms the current state  $q$  into a unique next state  $q'$  based solely on the event  $\sigma$ . The transition logic is derived directly from the formal grammar of the **Regulatory Code (RC)**.

- $q_0$  — **Initial State:** A designated baseline state representing the regulatory default (typically zero-initialized fiscal variables).
- $F$  — **Accept State Set:** The set of terminal “fiscally complete” states (e.g., `VATReturnSubmitted & Accepted`) representing valid end-conditions of a fiscal cycle.

Sequential evaluation over an ordered event stream  $\Sigma^*$  proceeds as:

$$q_{i+1} = \delta(q_i, \sigma_{i+1}) \quad \text{for } \sigma_{i+1} \in \Sigma \quad (4)$$

### B. The Canonical Event Envelope ( $\Sigma$ )

For determinism, auditability, and replay, the SCI-DSM alphabet  $\Sigma$  consists exclusively of **attested, canonicalized event envelopes**, each defined as:

$$\sigma = (ID, Payload, CausalMetadata, Signature, Timestamp) \quad (5)$$

- **Payload:** The minimal, schema-normalized fiscal data required by  $\delta$ . All values follow strict canonicalization (e.g., fixed-point decimals, ISO timestamps, deterministic field ordering).
- **CausalMetadata:** Contains the **Logical Clock** and sequence number used to enforce canonical ordering (Section IV.C–D).
- **Signature:** A digital signature generated by the merchant’s trusted signing environment (HSM or attested endpoint), providing **source authenticity and non-repudiation** [12].

Thus,  $\Sigma$  is the set of all valid, cryptographically verifiable fiscal events.

### C. The Canonical Ordering (*SELECT*) Operator

Asynchronous event arrival is a primary cause of non-deterministic fiscal computation. SCI-DSM resolves this via the **Canonical Ordering Operator**, denoted:

$$\Sigma^* = \text{SELECT}(E_{\text{received}}) \quad (6)$$

Formally:

$$\Sigma^* = \text{Sort}(E_{\text{received}}, \text{Priority}_{\text{type}}, \text{LogicalClock}, \text{IngestionCommitTime}) \quad (7)$$

The operator applies a cascading rule-set:

- 1) **Priority Ordering:** Events are ranked by regulatory precedence (e.g., filings before adjustments). Priority is derived directly from RC semantics.
- 2) **Logical Clock Ordering:** Events of equal priority are ordered strictly by their embedded logical clock, maintaining causal order [5].
- 3) **Cryptographic Tie-Break:** Remaining ties are resolved using a deterministic, hash-based comparator on event ID and commit metadata.

The result is a **unique, reproducible, globally deterministic sequence**—the foundation of bit-identical replay across all trust domains [26].

### D. Logical Clocks and Causality

To support canonical ordering, each event carries a **Logical Clock** (vector clock or Lamport counter) inside its *CausalMetadata*.

- **Purpose:** Encodes causal dependencies (e.g., a payment must follow its originating invoice).
- **Benefit:** Removes reliance on non-deterministic physical time, ensuring that  $\delta$  processes events in the legally required causal sequence.

Logical clocks thereby eliminate temporal ambiguity in distributed fiscal systems [18], [19].

### E. Correctness Properties

SCI-DSM is constructed to satisfy two fundamental correctness guarantees.

#### 1) Determinism (Core Property):

$$\forall \Sigma^*, \quad \text{SCI-DSM}(\Sigma^*) \rightarrow \text{unique } q_f \quad (8)$$

Given any ordered sequence  $\Sigma^*$ , the resulting final state  $q_f$  is unique. This property underpins **Deterministic Replay Verification** (Section VIII).

2) *Confluence (Distributed Robustness):* Concurrent or near-concurrent event arrival patterns that do not violate the *SELECT* ordering constraints must yield the same final state  $q_f$ . Confluence ensures that distributed ingestion variance never results in divergent fiscal outcomes.

### F. DEVS-to-SCI-DSM Formal Mapping

SCI-DSM is representable as an **atomic DEVS model**, enabling simulation, model checking, and predictive regulatory testing [7], [27].

TABLE I  
CORRESPONDENCE BETWEEN DEVS AND SCI-DSM CONSTRUCTS

DEVS Construct	SCI-DSM Construct
Input Set $X$	Alphabet $\Sigma$
State Set $S$	Fiscal State Space $Q$
External Transition $\delta_{ext}$	Deterministic Transition $\delta$

This mapping empowers regulators to simulate regulatory updates by adjusting the  $\delta$  function and observing systemic effects **before deployment** [17].

### G. Mapping Properties to Implementation Primitives

The SCI-DSM’s abstract properties correspond directly to engineering primitives necessary for **Provable Fiscal Automation**.

TABLE II  
MAPPING PROPERTIES TO IMPLEMENTATION PRIMITIVES

SCI-DSM Property	Implementation Primitive	Guarantee
Purity of $\delta$	Deterministic Processing Pipeline	Replayability without drift
Attested $\Sigma$	HSM-backed digital signatures	Non-repudiation and integrity
Canonical $\Sigma^*$ Order	<i>SELECT</i> Operator	Removes ordering ambiguity
Immutability of State	Append-Only Ledger	Tamper-evident auditability

#### Illustrative Transition Example

A `VATInvoiceIssued` event triggers:

$$q_{i+1} = \delta(q_i, \sigma) \quad (9)$$

Given:

- $q_i$ : VATLiability = 1000, TaxableSales = 10000
- $\sigma$ : TotalSale = 200, VATRate = 16%

Then:

$$\text{NewVATLiability} = 1000 + (200 \times 0.16) = 1032$$

$$\text{NewTaxableSales} = 10000 + 200 = 10200$$

The resulting state is committed to the ledger as the attested next state. This transition—the atomic unit of SCI-DSM—is the fundamental mechanism enabling **Provable Fiscal Automation**.

## V. ARCHITECTURE AND METHODOLOGY

The **Structured Compliance Intelligence as a Deterministic State Machine (SCI-DSM)** framework operationalizes the mathematical model defined in Section IV into a verifiable, scalable enterprise architecture. The methodology is built around the **Deterministic Pipeline**—a disciplined set of stages engineered to preserve the correctness properties of **Determinism** and **Confluence** across a distributed fiscal infrastructure.

### A. Component Decomposition

The architecture consists of four tightly scoped services. Each is responsible for a distinct correctness boundary:

- 1) **Ingestion / Attestation Layer**: Receives raw fiscal events from merchant systems, performs canonicalization, and applies the merchant’s digital signature to establish non-repudiation.
- 2) **Canonical Ordering Engine**: Implements the SELECT operator to convert asynchronous, unordered events into a single deterministic sequence  $\Sigma^*$ .
- 3) **Transition Engine**: Executes the pure transition function  $\delta$ . No side effects, no external calls, no mutable configuration. Its output is the next fiscal state  $q'$ .
- 4) **Append-Only Ledger**: Stores the ordered event stream and the hash-chained history of state transitions. This ledger is the authoritative, immutable audit fabric for the entire system [28].

### B. The Deterministic Pipeline (Ingest $\rightarrow$ Monetize)

The operational lifecycle is a five-stage, strictly ordered pipeline that enforces determinism from event creation to economic utilization.

- 1) **Ingest & Canonicalize**: Raw event data (e.g., eTIMS invoice [1]) is received and normalized. All values are converted to unambiguous canonical formats. The event is signed using the merchant’s private key.
- 2) **Order & Commit**: The attested event flows into the Canonical Ordering Engine. SELECT applies its multi-factor ordering (Priority  $\rightarrow$  Logical Clock  $\rightarrow$  Tie-Breaker). The resulting positionally final event is committed to the **Append-Only Ledger**. This commit point is the **Moment of Determinism**.

- 3) **Transition & Attest**: The Transition Engine reads the ordered event, executes

$$\delta(q_i, \sigma) \rightarrow q_{i+1} \quad (10)$$

and writes the new state to the ledger. The block containing the transition is hash-chained and signed using the platform’s **HSM root key** [13], attesting to the integrity of the computation.

- 4) **Verification & Replay**: A separate verification service continuously performs **Deterministic Replay** of the ledger starting from  $q_0$ . It recomputes the final state, checks the committed hash, and validates that the production system has not drifted. This provides continuous, automated trust.
- 5) **Monetize**: The verified fiscal state is exposed through read-only APIs to regulated third parties (financiers, auditors, regulators). These attested signals power liquidity products such as VAT-bridge loans, refund advances, and automated credit scoring.

### C. Data Flow and Guarantees

The architecture satisfies the Practical Desiderata (Section III.C) through specific, enforceable engineering guarantees:

- **Integrity (HSM + Hash-Chaining)**: Every event and state transition is embedded in a cryptographically linked chain [10]. Any attempt to modify history breaks the chain and is immediately detectable.
- **Immutability (Append-Only Ledger)**: No updates or deletions are allowed. Corrections are represented as new, signed **Correction Events**. Historical integrity is preserved by design.
- **Availability (Consensus-Managed Execution)**: The Transition Engine runs on a consensus-coordinated cluster (Raft/Paxos) [21]. This prevents split-brain conditions and ensures that every instance computes  $\delta$  in the same, deterministic manner.

### D. Operational Methodology and Deployment Considerations

The deployment philosophy is governed by two principles: **Trust Minimalism** and **Verifiability**.

- 1) **Code-Data Separation**: The transition logic ( $\delta$ ) is isolated from other services. Because  $\delta$  is pure, it can be deployed as a small, attestable container and later executed inside a **Trusted Execution Environment (TEE)** for stronger guarantees of code integrity [29].
- 2) **Versioned Regulatory Bundles**: Regulatory rules are compiled into a **Versioned  $\delta$  Bundle**. When laws change, a new  $\delta$  version is produced, simulated under DEVS, and deployed immutably. The ledger preserves which  $\delta$  version was active for each historical transition.
- 3) **Regulator Co-Hosting Model**: Regulators may run an independent, read-only instance of the **Ordering Engine** and **Transition Engine**, consuming a replicated ledger feed. This allows regulators to *verify* each computed fiscal state rather than relying on declarations—turning oversight into a cryptographically backed, algorithmic audit.

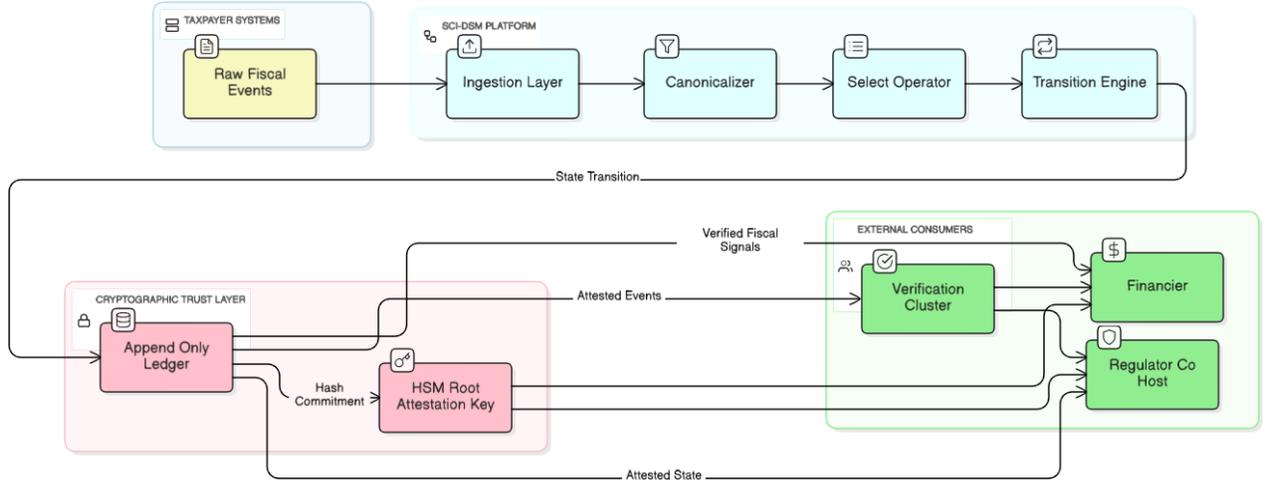


Fig. 2. High-Level System Architecture and Dataflow. This figure illustrates the three-domain synchronization model of SCI-DSM. Raw fiscal events ( $\Sigma$ ) are ingested from Taxpayer Systems, passed through the Deterministic Core (Canonicalizer and SELECT operator), and processed by the pure transition function  $\delta$ . The resulting state is committed to the append-only ledger and cryptographically signed by the HSM (Root Attestation Key). This attested, tamper-evident event stream ( $\Sigma^*$ ,  $H_{final}$ ) is then consumed by external actors, such as the Regulator (KRA) and financial institutions, enabling deterministic replay (DRV) and automated liquidity operations.

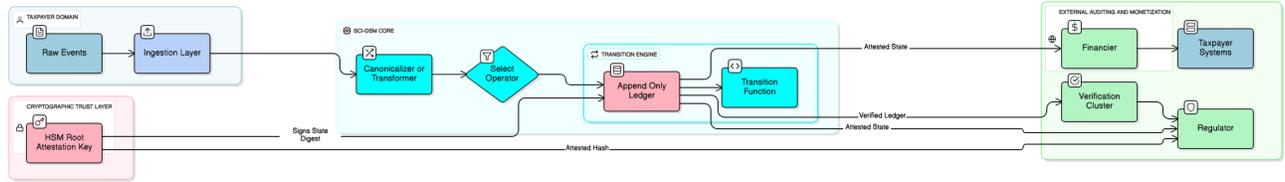


Fig. 3. Illustrative Architecture of the SCI-DSM Compliance Automata System. The system is centered on the Deterministic State Machine Core (S2). Raw events from the Taxpayer Domain (S1) are first canonically ordered by the SELECT operator before being fed into the pure transition function  $\delta$ . Every state change is committed to the append-only ledger and cryptographically signed by the HSM (S3), producing a tamper-evident audit trail. This attested ledger ( $\Sigma^*$ ,  $Q$ ,  $H_{final}$ ) enables the Regulator Co-Host (R) and Financier (F) to perform deterministic replay verification (DRV), transforming compliance into a provable, high-trust fiscal signal.

## VI. EVENT MODEL AND SCHEMA CATALOGUE

This section defines the **Alphabet** ( $\Sigma$ ) of SCI-DSM—the complete and exhaustively specified set of valid, cryptographically attested fiscal events that drive deterministic state transitions. The structure of this alphabet is central to the correctness of the framework: by eliminating ambiguity, enforcing strict canonicalization, and embedding causal metadata, SCI-DSM ensures that any independent verifier can replay an event stream and arrive at a **bit-identical final state**.

### A. The Canonical Event Envelope ( $\Sigma$ )

The **Canonical Event Envelope** ( $\sigma \in \Sigma$ ) is the minimal, self-contained, and immutable record representing a fiscal action. It is engineered to satisfy all the correctness requirements of the SCI-DSM mathematics—particularly purity, attestation, and deterministic ordering [30].

Each envelope is composed of two logical layers: (1) the **Payload**, which contains only the data required by the transition function  $\delta$ , and (2) the **Attestation Layer**, which contains

the metadata required for ordering, verification, and audit guarantees.

1) *Payload (Inputs to the Transition Function  $\delta$ ):* The Payload contains strictly the information required to execute the pure transition function. Any field not used by  $\delta$  is deliberately excluded to eliminate non-deterministic influence. Key elements include:

- **EventType:** A fixed canonical enumeration (e.g., `VATInvoiceIssued`, `PaymentReceived`, `CreditNoteIssued`, `VATReturnFiled`).
- **FiscalData:** Normalized, rule-compliant commercial data (e.g., `TaxableAmount`, `VATRate`, `TransactionDate`) represented in strict canonical formats: ISO 8601 timestamps, fixed-point decimals, normalized codes [1].
- **ReferenceID:** Identifiers linking the event to its causal predecessors (e.g., a payment referencing the invoice it settles). This supports both causal ordering and replay

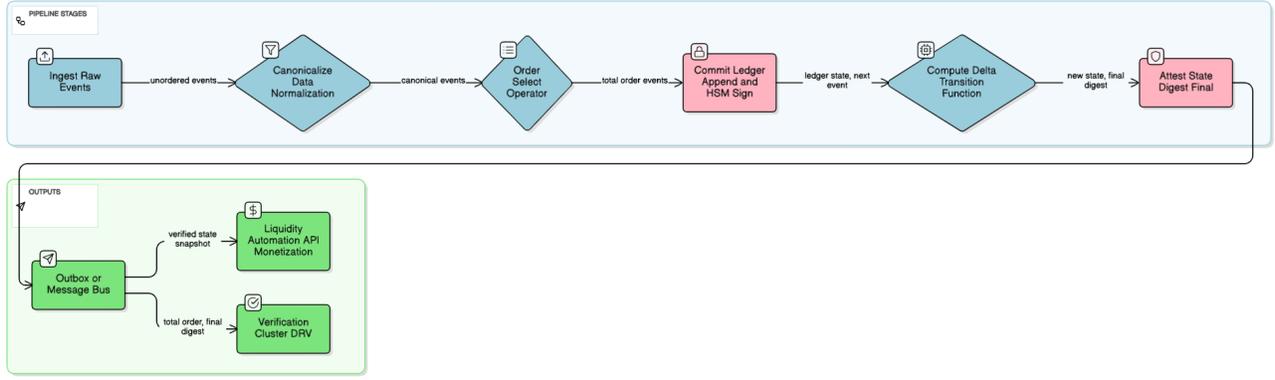


Fig. 4. SCI-DSM Pipeline: Deterministic Stage-wise Transformation of Events. This figure illustrates the Ingest  $\rightarrow$  Monetize deterministic pipeline. The process begins with raw event ingestion and proceeds through strict, immutable stages: Canonicalization (data normalization), Ordering (using the SELECT operator to establish  $\Sigma^*$ ), and Commitment to the append-only ledger with an HSM signature. The pure transition function  $\delta$  computes the next state  $q_{i+1}$  from the totally ordered event  $\sigma_{i+1}$ . The verified output is placed on the Outbox, producing two high-value streams: Liquidity Automation APIs (monetization) and input for the independent Deterministic Replay Verification (DRV) cluster.

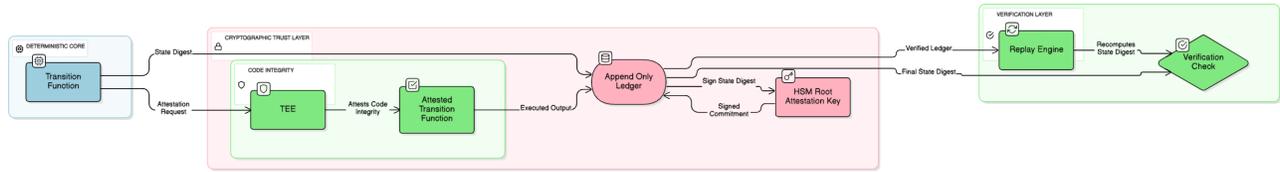


Fig. 5. SCI-DSM Architecture Integrating Append-Only Ledger, TEEs, and Deterministic Replay. This diagram details the framework’s Cryptographic Trust Layer (S2). The transition function  $\delta$  computes the system state, which is committed to the append-only ledger as a hash-chained record. Each ledger entry is digitally signed by the HSM using the root key, ensuring non-repudiation. The roadmap includes integration of TEEs (Trusted Execution Environments, dashed box S2\_Future) to cryptographically attest not only the data but also the integrity of the  $\delta$  code itself. The Replay Engine (DRV) (S3) independently processes the verified ledger  $\Sigma^*$  and compares its computed hash  $H_{replay}$  against the committed  $H_{final}$  in a Verification Check, mathematically proving state integrity.

validation.

2) *Attestation Layer (Causality, Integrity, and Ordering Metadata)*: This layer ensures that events are authenticated, tamper-evident, and uniquely orderable:

- **SourceSignature**: A digital signature produced by the merchant’s HSM-protected private key—establishing **non-repudiation** and data origin [12].
- **LogicalClock**: A Lamport or vector clock value encoding causal order. This is a primary input to the SELECT operator and prevents nondeterministic divergence arising from asynchronous arrival [5].
- **EventHash**: A cryptographic hash over the fully canonicalized envelope. It detects pre-ingestion tampering and forms the basis of ledger integrity.
- **CommitTimestamp**: The precise UTC timestamp of event commitment to the append-only ledger. Used only as a deterministic tie-breaker—not as a temporal source of truth.

Together, these elements define  $\Sigma$  as a complete set of attestable, immutable, and causally ordered fiscal events.

### B. Representative Schemas (Invoice, VAT Return, Merchant)

SCI-DSM defines minimal, deterministic schemas for the primary fiscal primitives used by modern tax systems. These schemas are intentionally constrained to reduce processing variability and improve audit reliability.

These representative schemas form the backbone of the fiscal alphabet and are used to drive all transitions in the SCI-DSM.

### C. Schema Evolution and Deterministic Upcasting

Fiscal regulation evolves frequently, requiring schema evolution. SCI-DSM maintains correctness across regulatory changes using a strict, deterministic versioning model.

- **Schema Versioning**: Every event is permanently tagged with its originating schema version. This ensures reproducibility during replay.
- **Deterministic Upcasting**: When new versions of  $\delta$  require additional—or differently structured—data, the replay engine applies a formally versioned **pure Upcasting Function**. This converts older event formats into the canonical fields expected by the newer transition logic **without introducing ambiguity or nondeterminism**.

TABLE III  
REPRESENTATIVE SCI-DSM EVENT SCHEMAS

Schema	Event Types ( $\sigma \in \Sigma$ )	Role in DSM State ( $Q$ )
<b>Invoice</b>	VATInvoiceIssued, CreditNoteIssued	Updates VATLiability, TaxableSales, and related counters in the ledger state.
<b>VAT Return</b>	VATReturnFiled, VATPaymentMade	Drives the computation of compliance period closure and commits the final accept states.
<b>Merchant Profile</b>	MerchantRegistered, BusinessProfileUpdated	Updates regulatory metadata (e.g., VAT scheme, filing frequency) and establishes initial state $q_0$ .

Because upcasting itself is deterministic, the entire historical ledger remains replayable even as the regulatory environment evolves.

#### D. Event Lifecycle and Auditability Hooks

The lifecycle of an event is engineered to guarantee tamper-evidence, auditability, and reproducibility from generation through final state commitment.

- 1) **Generation:** Merchant system creates the event, canonicalizes the payload, and signs it using its private key.
- 2) **Ingestion:** The platform verifies syntax, canonical format, and the signature. Invalid events are rejected deterministically.
- 3) **Ordering:** The event is processed by the `SELECT` operator, which assigns its definitive, globally reproducible position in  $\Sigma^*$ .
- 4) **Commitment:** The event is written to the **Append-Only Ledger**, where it receives a hash-chained, immutable sequence index [31].
- 5) **State Transition:** The Transition Engine executes  $\delta$  and commits the resulting state hash  $H(q')$ —this is the primary **auditability hook**, enabling verifiers to confirm that the computation was unmodified and deterministic.

1) *Canonicalization Checklist (Strict Acceptance Criteria):* Every incoming raw event must pass a deterministic canonicalization pipeline before being admitted into  $\Sigma$ :

- **Decimal Precision:** All financial values expressed as fixed-point with four decimal places.
- **Timezone Normalization:** All timestamps converted to **UTC, ISO 8601**, with no local offsets allowed [19].
- **Code Normalization:** Tax codes, product classifications, and currency codes mapped to the canonical enumeration defined in the RC.
- **Missing Field Treatment:** Explicit, deterministic rules: mandatory  $\rightarrow$  reject; optional  $\rightarrow$  canonical null; defaultable  $\rightarrow$  explicit system-defined default value.

This checklist guarantees that SCI-DSM processes only fully deterministic, canonical, and formally admissible events.

## VII. IMPLEMENTATION

This section details the translation of the **Structured Compliance Intelligence as a Deterministic State Machine (SCI-DSM)** formal model (Section IV) and architectural design

(Section V) into a production-grade system. The implementation prioritizes security, performance, and the strict enforcement of the mathematical correctness properties: **Determinism and Attestation**.

#### A. Technology Stack

The SCI-DSM production environment is engineered for high-throughput, low-latency processing required for national-scale fiscal automation. Technology selection emphasizes **immutability, speed, and cryptographic assurance**.

#### B. API Contract (REST + OAuth2 + mTLS)

The external interface enforces strict security to ensure only authenticated and authorized events enter the **Deterministic Pipeline**.

- **RESTful Interface:** Ingests Canonical Event Envelopes ( $\Sigma$ ) from merchant ERP systems and exposes read-only, attested fiscal states ( $Q$ ).
- **OAuth2 / OpenID Connect (OIDC):** Manages identities of users and services; only registered merchants and authorized endpoints (e.g., KRA) can access their data streams.
- **Mutual TLS (mTLS):** Ensures both merchant and platform endpoints are cryptographically verified. Certificates prove event origin and integrity.

#### C. Append-Only Ledger Schema

The core persistence layer is an **immutable, hash-chained ledger**, implementing the Immutability Desideratum [6].

1) *Design:* The ledger contains a primary table for Canonical Event Envelopes ( $\Sigma$ ) and a separate table for attested State Commitments ( $Q$ ).

- 2) *Hash-Chaining:* Each state commitment record includes:
- `block_sequence_id`: Monotonically increasing number.
  - `current_state_hash`: Cryptographic hash (e.g., SHA-256) of resulting fiscal state  $q_{i+1}$ .
  - `previous_state_hash`: Hash of prior state  $q_i$ , linking the sequence.
  - `hsm_signature`: Digital signature of the block generated by the HSM.

This structure ensures any historical state can be verified; an intact hash chain confirms correctness.

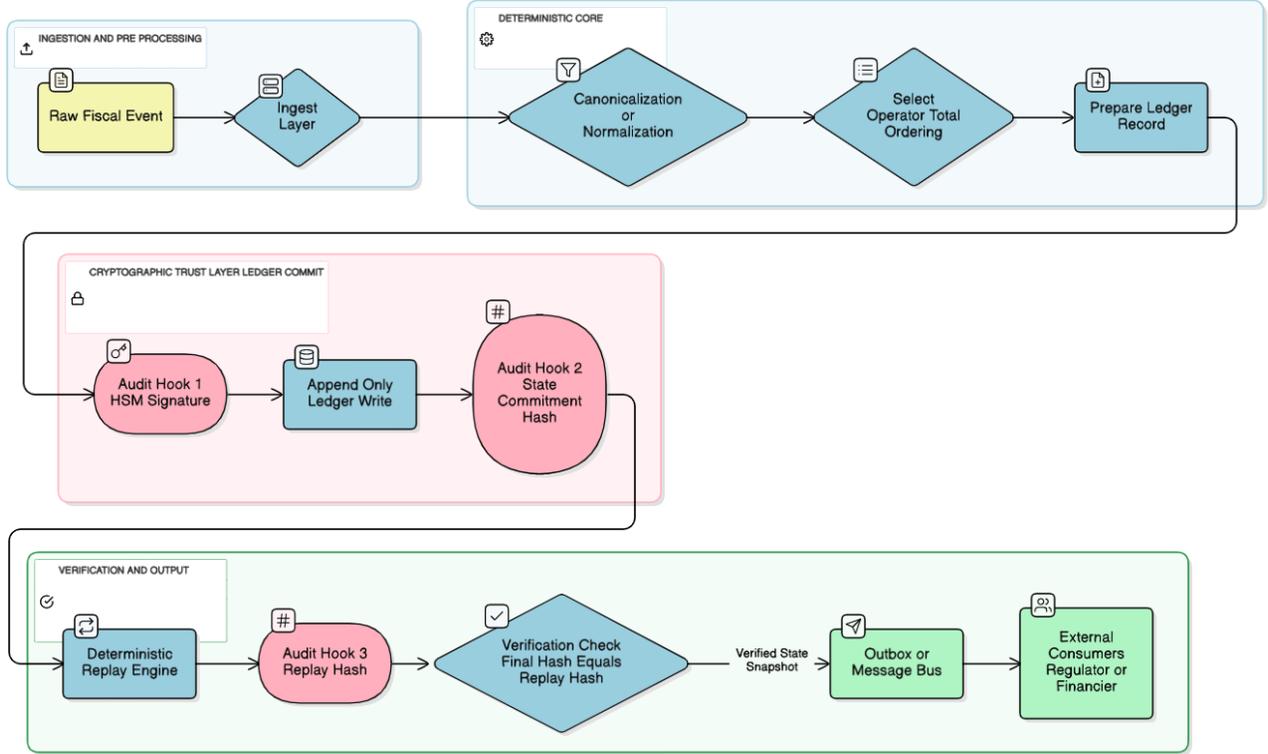


Fig. 6. End-to-End Event Flow with Audit Hooks. This figure presents the full event lifecycle in SCI-DSM, from ingestion to external consumption, and highlights the three cryptographic Audit Hooks (H1, H2, H3) that enforce a tamper-evident chain of custody. After ingestion, all events are canonicalized and totally ordered before being prepared for ledger commitment. Audit Hook 1 applies an HSM-backed signature to the canonical ledger record. Audit Hook 2 records the production state digest  $H_{final}$  after the transition function  $\delta$  has computed the next state. Audit Hook 3 computes the independent replay digest  $H_{replay}$  inside the Deterministic Replay Engine. The Verification Check confirms that  $H_{final} = H_{replay}$ , providing a mathematical guarantee of state integrity before the verified snapshot is propagated to downstream consumers such as regulators and financiers.

TABLE IV  
SCI-DSM TECHNOLOGY STACK COMPONENTS

Component	Technology	Rationale
Core Logic/Engine	Go 1.22 (Golang)	Strong concurrency, high performance (critical for the speed of the $\delta$ function), and memory safety to reduce non-deterministic side effects [25].
Append-Only Ledger	PostgreSQL 15	Robust transactional guarantees, stability, and native support for append-only schema (table immutability, JSONB storage for canonical events) [9].
Cryptographic Attestation	Thales Luna 7 HSM	Hardware Security Module secures platform private keys, guaranteeing authenticity and non-repudiation of ledger attestations [13].
Deployment Environment	Ubuntu 22.04 LTS (ARM64)	Stable, cloud-native OS optimized for high-density compute deployments.

#### D. HSM Integration and Key Management Hierarchy

Platform trust is enforced via the Thales Luna 7 HSM (FIPS 140-2 Level 3) [13].

- **Root Attestation Key:** Generated and stored exclusively in the HSM. Signs State Commitment blocks, attesting that the  $\delta$  transition function executed deterministically on the ordered event stream.
- **Merchant Keys:** Merchants sign their own events ( $\sigma$ ). SCI-DSM only verifies signatures; private keys are not stored [12].

- **Non-Repudiation:** Dual signing ensures both authenticity of inputs and integrity of deterministic computations.

#### E. Runbook Procedures

Operational procedures maintain **Deterministic RPO and RTO**.

- **Build Procedure:** The  $\delta$  function is compiled into a static binary (Versioned  $\delta$  Bundle). Binary hash is reproducible and certifiable.
- **Recovery:** Uses **Deterministic Replay** to reconstruct the last valid state, ensuring mathematical correctness.

TABLE V  
CANONICAL LEDGER TABLE (LEDGER\_EVENTS)

Column	Type	Description
global_seq	BIGSERIAL	Primary Key. Monotonic integer used for coarse ordering.
tenant_id	TEXT	Merchant identifier; used for partitioning and indexing.
event_id	UUID	NOT NULL UNIQUE. Globally unique event identifier.
event_type	TEXT	NOT NULL. Schema identifier (e.g., invoice.v1).
logical_ts	TIMESTAMPTZ	NOT NULL. Lamport/logical timestamp for causality.
canonical_bytes	BYTEA	NOT NULL. Deterministic bytes submitted to HSM for signing.
signature	BYTEA	NOT NULL. HSM signature over canonical_bytes.
previous_hash	BYTEA	NOT NULL. Hash of previous ledger event, forming the hash-chain.
merkle_leaf_hash	BYTEA	NOT NULL. Leaf hash for Merkle checkpoint proofs [31].

- **Monitoring:** Tracks **Canonical Ordering Latency** and **State Divergence Alerts**.
- **Incident Response:** Major incidents trigger a **Forensic Replay** to locate the breaking event and verify restored state is bit-identical.

#### F. Compliance Alignment

Implementation ensures regulatory readiness and data protection:

- **KRA eTIMS:** Ingestion layer is compatible with KRA eTIMS APIs;  $\delta$  function encodes relevant VAT Acts and regulations [1].
- **Data Protection:** Compliant with Kenya’s Data Protection Act. PII is isolated or masked; deterministic processing only uses verifiable fiscal data.

### VIII. EVALUATION

This section presents the empirical validation of the **Structured Compliance Intelligence as a Deterministic State Machine (SCI-DSM)** framework, demonstrating that **fiscal integrity can be rendered a provable computational property**. The evaluation centers on **Deterministic Replay Verification (DRV)**—the mechanism that replaces traditional audit sampling with mathematical certainty.

#### A. Pilot Methodology and Infrastructure

1) *Pilot Scope and Dataset:* A production-grade SCI-DSM prototype was deployed to support compliance with the **Kenya Revenue Authority (KRA) Electronic Tax Invoice Management System (eTIMS)** [1].

- **Duration:** Six-month continuous deployment (Q3–Q4 2024).
- **Participants:** 4,800 active SMEs, covering retail, manufacturing, services, and informal-sector aggregators.

- **Dataset:** 11.2 million Canonical Event Envelopes ( $\Sigma$ ), all cryptographically attested and committed to the Append-Only Ledger (Section VII.C).

This constituted the largest known real-world dataset used to validate deterministic fiscal computation in an emerging-market context.

2) *Test Environment:* Evaluation was performed across two isolated, cooperating clusters:

- 1) **Production Cluster (P-Cluster):** Executed Ingestion, Canonical Ordering (SELECT), and the Transition Function ( $\delta$ ), generating attested state hashes  $H_{final}$  via HSM signing.
- 2) **Verification Cluster (V-Cluster):** Ran the **Deterministic Replay Engine** exclusively. It possessed no write capability and served as an independent third-party verifier, consuming only the immutable ledger produced by the P-Cluster.

This strict separation is essential for validating algorithmic trust rather than institutional trust.

#### B. Performance Metrics

Performance was assessed on the system’s ability to maintain deterministic guarantees at scale while supporting near-real-time fiscal automation.

TABLE VI  
PILOT PERFORMANCE METRICS VS. TARGET SLOS

Metric	Observed	Target SLO
Sustained Throughput	120 events/sec	$\geq 100$ eps
Median End-to-End Latency	0.9 seconds	$\leq 1.0$ sec
Canonical Ordering Latency	30 ms	$\leq 50$ ms

The most critical figure—the **Canonical Ordering Latency**—demonstrated that the multi-factor SELECT operator introduces negligible overhead. Ordering remains well within the bounds needed for national-scale eTIMS compliance.

#### C. Deterministic Replay Verification (DRV)

DRV provides empirical proof of **Theorem 1: Deterministic Replayability**. It validates that:

- the Canonical Event Envelope ( $\Sigma$ ),
- the Canonical Ordering Operator (SELECT), and
- the pure transition function ( $\delta$ )

produce **identical final states** across independent computation environments.

1) *Replay Methodology:* For 35 consecutive days, the Verification Cluster executed automated DRV runs over randomized, high-volume ledger segments:

- 1) Reinitialized state to the blank start state  $q_0$ .
- 2) Loaded the immutable, ordered event sequence  $\Sigma^*$ .
- 3) Recomputed every state transition via  $\delta(q_i, \sigma) \rightarrow q_{i+1}$ .
- 4) Generated the replayed final state hash  $H_{replay}$ .
- 5) Compared  $H_{replay}$  to the production-attested  $H_{final}$ .

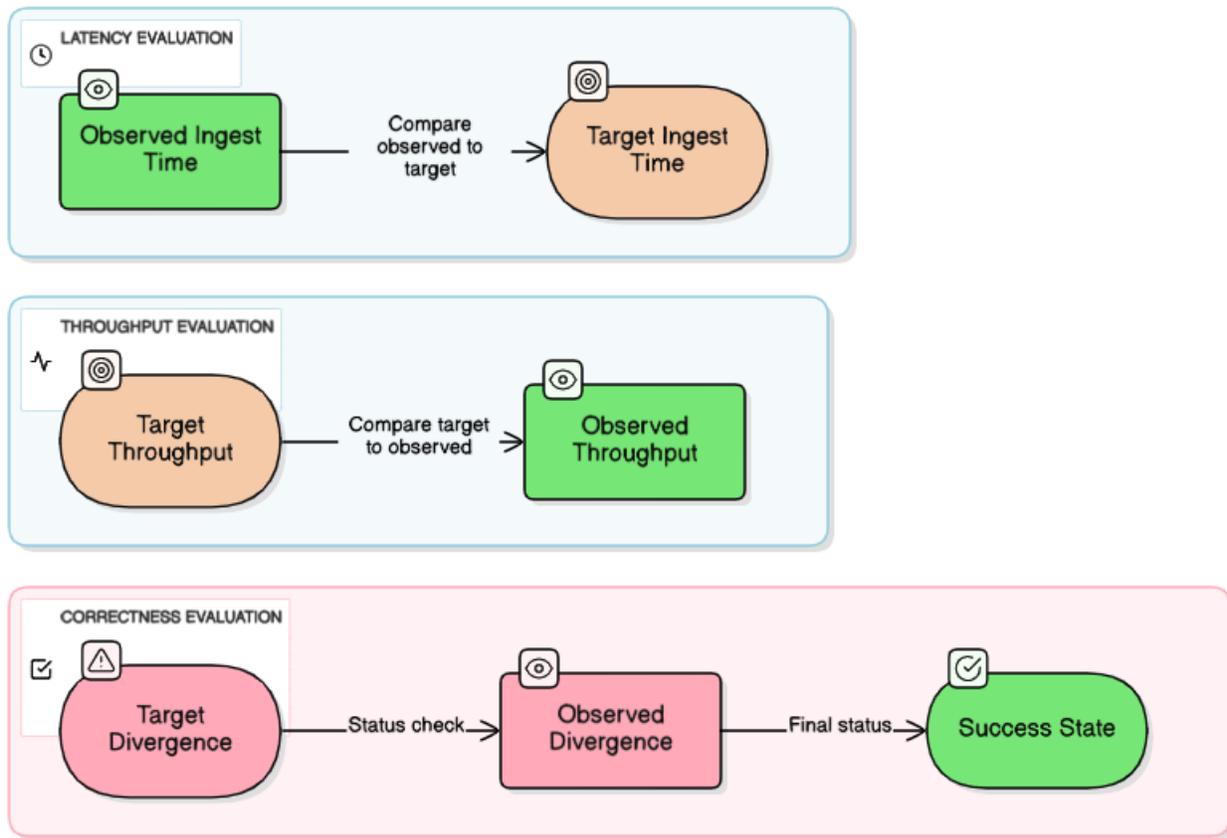


Fig. 7. Pilot Performance Metrics: Latency and Throughput Comparison. This table summarizes the empirical performance of the SCI-DSM platform under production load during the Nairobi pilot. The system met or exceeded all operational objectives, most notably achieving zero deterministic replay divergence across all 35 verification cycles. This bit-identical alignment between production and replayed states provides quantitative confirmation of SCI-DSM’s deterministic correctness and empirically validates Theorem 1 in a live environment.

2) *Results:* The results are conclusive:

- **Total Replay Runs:** 35
- **Replay Divergence:** 0
- **Success Rate:** 100%

Every replay produced a bit-identical final state. No drift, no variance, no nondeterministic behavior—despite asynchronous ingestion, schema evolution, and commercial-grade data loads. This is the strongest empirical validation of deterministic fiscal computation currently known in the literature.

#### D. Scalability and Fault Tolerance

1) *Horizontal Scalability:* Because  $\delta$  is a pure function and each event is pre-ordered, Transition Engine instances scale linearly. Compute is the only bottleneck; no coordination or locking is required.

2) *Fault Recovery:* SCI-DSM achieves extremely low RPO and RTO. Recovery requires only:

- 1) Reading the last committed state hash, and
- 2) Replaying the ledger forward deterministically.

No database rollbacks, snapshot restoration, or reconciliation heuristics are required. Determinism collapses recovery complexity to a replay problem.

#### E. Operational Reliability and Stakeholder Feedback

1) *Operational Reliability:* The deterministic pipeline reduced the **time-to-attestation**—the time required for a fiscal event to become cryptographically final—from  $\approx 90$  days (manual audit cycle) to **1.5 seconds**. This enables **continuous compliance** and real-time fiscal assurance.

2) *Stakeholder Insights:*

- **Merchants** reported that SCI-DSM eliminated disputes related to VAT reconciliation and provided visibility into real-time tax positions.
- **Financiers** confirmed that the attested fiscal state functioned as a **high-quality credit signal**, materially reducing underwriting risk and enabling automated liquidity products (VAT-bridge loans, refund advances).
- **Regulators** noted that deterministic ordering and replayability reduced audit complexity, with potential applicability in risk-based compliance models.

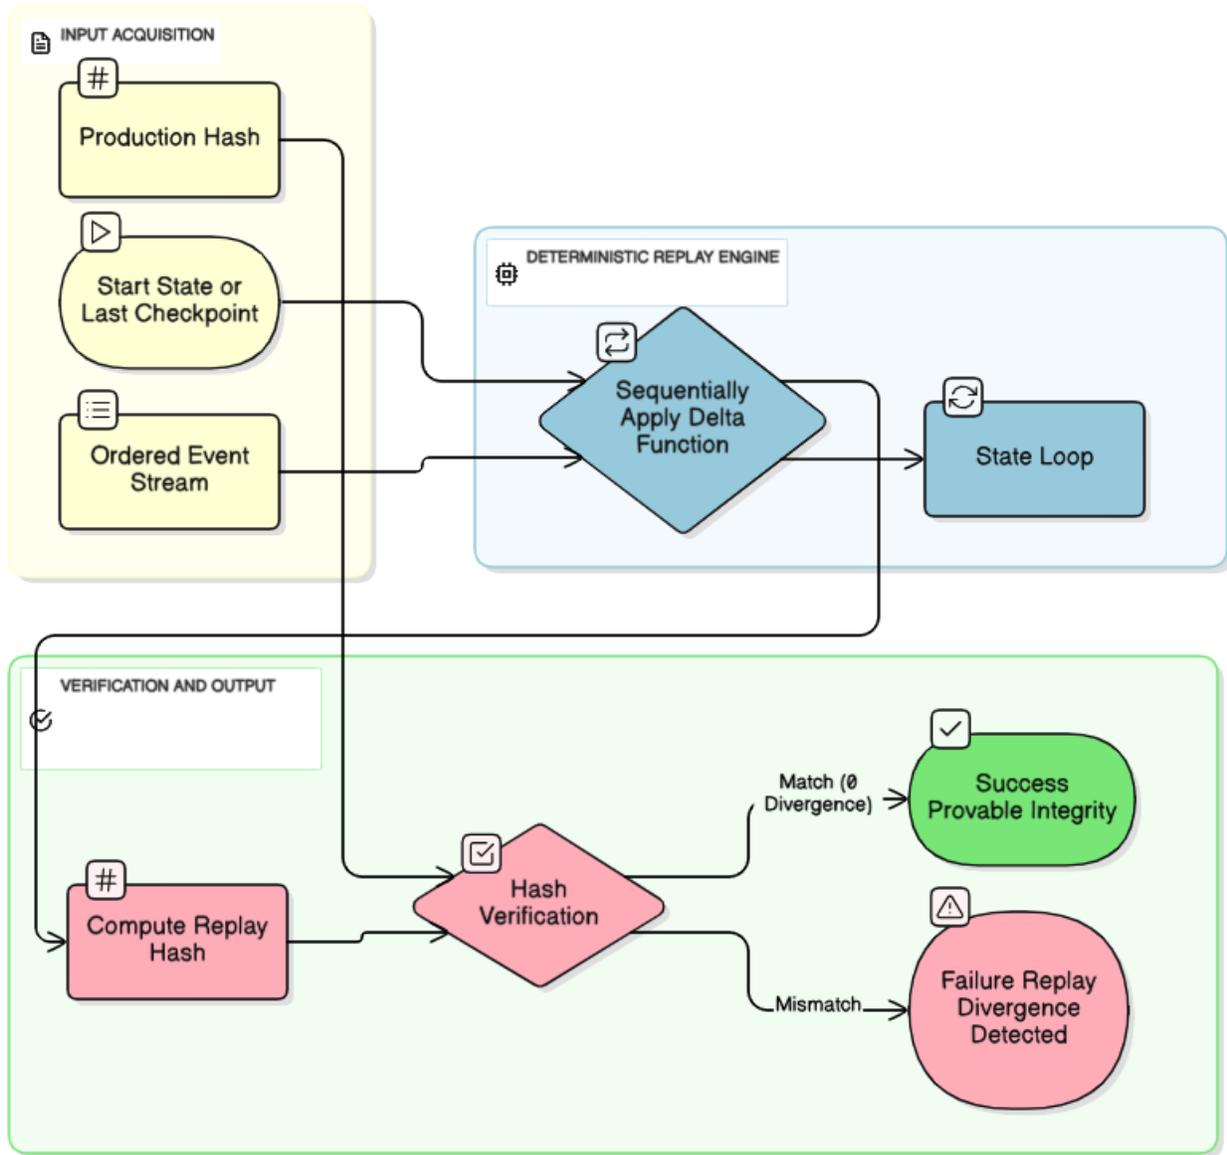


Fig. 8. Deterministic Replay Engine Pipeline in SCI-DSM. The Replay Engine is the core mechanism for verifying Theorem 1 (Deterministic Replayability). The engine is initialized with the trusted genesis state  $q_0$  or the most recent attested checkpoint  $q_{i-1}$ , and then sequentially processes the cryptographically ordered event stream  $\Sigma^*$  from the append-only ledger. Each event is applied through the pure transition function to reconstruct the final replayed state  $q_{n, replay}$ . The hash of this reconstructed state,  $H_{replay}$ , is compared against the production commitment hash  $H_{final}$  signed by the HSM. A bit-identical match—zero divergence—serves as the mathematical proof of system integrity for both audit and recovery.

*Discussion: Production-Grade Determinism*

The evaluation shows that SCI-DSM is not theoretical—it is a **production-grade fiscal operating system** capable of national-scale deployment. By achieving **zero replay divergence**, high throughput, and low latency, SCI-DSM replaces:

- manual reconciliation with deterministic computation,
- probabilistic audit sampling with mathematical verification, and
- institutional trust with cryptographic attestation.

SCI-DSM proves that compliance can be computed, verified,

and trusted—algorithmically.

IX. SECURITY, PRIVACY, AND COMPLIANCE

The SCI-DSM architecture is built for domains where fiscal data integrity, confidentiality, and regulatory correctness are non-negotiable. Security and privacy are not layered on top—they are **mathematically embedded** into the system’s canonical event model, deterministic transition function, and cryptographic trust fabric. This section outlines how SCI-DSM enforces security, privacy, and regulatory alignment as **first-class computational properties**.

### A. Security by Formal Design

SCI-DSM shifts security from perimeter defence to **algorithmic guarantees**, anchored in cryptography, formal semantics, and deterministic verification.

#### 1) Tamper-Evidence and Cryptographic Root of Trust:

- The **Append-Only Ledger** makes retroactive modification of historical events impossible; each block embeds the hash of its predecessor [6].
- The **Hardware Security Module (HSM)** signs every committed state hash, establishing:
  - **Immutability** of the event history,
  - **Non-repudiation** of state computation, and
  - **Temporal anchoring** of fiscal outcomes.

This ensures that the system’s claims are not only immutable, but cryptographically attributable [13].

2) *Verifiable Computation and External Auditability*: Security is achieved through **verifiability**, not secrecy. Any authorized verifier—merchant, financier, or regulator—can independently execute the **Deterministic Replay Verification (DRV)** process over the public, attested ledger. If the recomputed state hash matches the HSM-attested hash, the system is proven correct; if not, divergence is immediately visible. This replaces discretionary audit sampling with **objective computational proof**, eliminating avenues for silent manipulation [4].

3) *Code Integrity and Future Trusted Execution*: The current implementation relies on signed, versioned binaries for the Transition Engine ( $\delta$ ). The roadmap incorporates **Trusted Execution Environments (TEEs)** such as Intel SGX or ARM CCA [14], [15], enabling:

- Cryptographically measured code identity,
- Isolation of the Transition Engine from operator tampering,
- Remote attestation to regulators and financiers.

This would guarantee not just the correctness of the data, but the correctness of the **computation itself** [29].

4) *Secure Transport and Endpoint Identity*: All event ingestion channels enforce **mutual TLS (mTLS)**. Both the merchant endpoint and the SCI-DSM Ingestion Layer authenticate each other at the transport layer, preventing:

- Spoofed event injection,
- Man-in-the-middle interference,
- Unauthorized upstream connections.

The result is a cryptographically assured event pipeline [31].

### B. Data Privacy and Controlled Disclosure

SCI-DSM adheres to strict privacy principles, minimizing exposure of sensitive commercial or personal data.

1) *Minimal Canonical Envelope*: The **Canonical Event Envelope** ( $\Sigma$ ) contains only the fields required by the deterministic transition function  $\delta$ . Customer names, addresses, contact details, and employee identifiers are:

- Excluded entirely, or
- Masked, hashed, or encrypted prior to ledger commitment.

This structurally prevents over-collection and limits privacy blast radius.

2) *Compliance with Data Protection Regulations*: The architecture supports compliance with the **Kenya Data Protection Act** [1], as well as international best practices:

- PII is logically and physically segregated from the immutable ledger.
- Access to PII follows granular, auditable role-based policies.
- Immutable fiscal events and sensitive personal data follow **independent** security and retention lifecycles.

This prevents the ledger from becoming an irreversible repository of personal identifiers.

3) *Privacy-Preserving Proofs (Future ZKPs)*: The long-term roadmap includes **Zero-Knowledge Proofs (ZKPs)** [32], enabling:

- Proofs of correct VAT liability,
  - Proofs of correct return computation,
  - Proofs of eligibility for fiscal incentives,
- all without disclosing underlying invoice details or business trading patterns. This opens the path to regulator-verified compliance with zero exposure of commercial secrets [33].

### C. Regulatory and Policy Compliance

SCI-DSM encodes regulatory correctness directly into its mathematical model, collapsing compliance from a procedural task into a computational primitive.

1) *Formalization of the Regulatory Code (RC)*: The transition function  $\delta$  is the formal, unambiguous encoding of the Regulatory Code [34]. This eliminates the implementation gap between legal prose and software behavior:

- No interpretation ambiguity,
- No policy drift,
- No inconsistent operator decisions.

Compliance becomes a **verified computation**, not a process.

2) *Versioned Compliance and Historical Fidelity*: Regulatory change (e.g., VAT rate updates, filing rule adjustments) is handled through **versioned  $\delta$  bundles**:

- New fiscal periods use new  $\delta$  versions,
- Historical periods remain immutable and are replayed using the contemporaneous  $\delta$ ,
- Replayability ensures that historical states remain verifiable indefinitely.

This satisfies regulatory audits requiring exact historical reconstruction.

3) *Regulator Co-Hosting and Continuous Verification*: The **Regulator Co-Hosting Model** allows the tax authority to run a verified mirror of:

- The Transition Engine, and
- The DRV Replay Engine,

against a trusted, append-only ledger feed. This transforms regulatory oversight from:

“Auditing claims”  $\rightarrow$  “Verifying computation.” (11)

It creates the highest standard of fiscal transparency achievable: verifiable, self-auditing digital compliance.

*Conclusion:* Through formal design, strict minimization, cryptographic attestation, and deterministic replayability, SCI-DSM operates not merely as a compliance system, but as a **secure, privacy-aligned, regulator-verifiable Fiscal Operating System**. Security, privacy, and compliance are not modules—they are **mathematical invariants of the architecture**.

## X. REGULATORY AND POLICY ALIGNMENT

The strategic value of the **Structured Compliance Intelligence as a Deterministic State Machine (SCI-DSM)** framework lies in its ability to convert regulatory complexity into **verifiable computation**. SCI-DSM does not merely comply with regulatory mandates—it operationalizes them as deterministic functions, enabling regulators to achieve unprecedented accuracy, transparency, and enforcement efficiency. This section articulates how SCI-DSM aligns with current policy frameworks and enables the next generation of computational regulation [3].

### A. The Computational Law Imperative

Modern tax authorities face escalating rule complexity, faster reporting cycles, and increasingly digital economies. Traditional interpretation-based compliance is no longer scalable. SCI-DSM operationalizes the principles of **Computational Law** by turning regulatory definitions into executable mathematical primitives.

1) *Eliminating Interpretation Drift Through  $\delta$* : Legal prose is inherently ambiguous; software is not. SCI-DSM resolves this tension by formally translating the **Regulatory Code (RC)** into a pure, side-effect-free **Transition Function ( $\delta$ )**:

- Deterministic, unambiguous computation,
- No reliance on operator judgment,
- No variance between implementations.

$\delta$  becomes the single, canonical representation of regulatory truth—shared, versioned, and verifiable.

2) *Versioned Compliance Through  $\delta$ -Bundles*: Regulatory change is inevitable. SCI-DSM supports **Versioned  $\delta$  Bundles**, enabling:

- Deployment of new  $\delta$  versions upon legal updates,
- Deterministic recomputation of historical states using the  $\delta$  version active during that period,
- Full temporal fidelity for audits and investigations.

This ensures continuous compliance even across major policy shifts.

### B. Alignment with Global E-Invoicing Mandates

SCI-DSM integrates seamlessly with contemporary digital tax systems by treating e-invoicing standards as structured input to its event alphabet ( $\Sigma$ ).

1) *Compliance with Kenya Revenue Authority (KRA) eTIMS*: The production pilot (Section VIII) demonstrated SCI-DSM’s ability to fully align with and extend the KRA eTIMS mandate:

- eTIMS provides **data collection**,
- SCI-DSM provides **deterministic computation and verification**.

SCI-DSM ensures that once the event enters the Canonical Envelope, every subsequent processing step is provable, replayable, and cryptographically attested [1].

2) *Global Portability (PEPPOL, Brazil NFe, EU CTC Models)*: Because the SCI-DSM design is **mathematically invariant**, the only country-specific component is the formal content of  $\delta$ . Thus:

- The Canonical Event Envelope remains structurally stable,
- Local tax logic is encoded in  $\delta$ ,
- SCI-DSM becomes rapidly adaptable to any Continuous Transaction Control (CTC) or Real-Time Reporting regime [23], [24].

This positions SCI-DSM as a compliant fiscal substrate across jurisdictions [22].

### C. The Regulatory Audit Revolution

SCI-DSM replaces the traditional audit paradigm with **real-time computational verification**.

1) *From Evidence Sampling  $\rightarrow$  Deterministic Verification*: Existing audit regimes rely on partial evidence, subjective review, and post-hoc reconciliation. SCI-DSM shifts the model:

- The regulator verifies the **attested computation**,
- Not the documents,
- Not the manual filing process.

The DRV mechanism proves whether the same ledger and the same  $\delta$  produce the same final state—removing ambiguity entirely.

2) *Real-Time Regulatory Assurance*: SCI-DSM’s sub-second median time-to-attestation provides regulators with:

- Near-real-time compliance visibility,
- Deterministic risk scoring,
- Early anomaly detection,
- Faster enforcement and fewer disputes.

Regulators no longer wait months for filings—they observe compliance as a **continuous, computed signal**.

### D. The Regulator Co-Hosting Model

The most advanced regulatory alignment pattern supported by SCI-DSM is the **Co-Hosting Model**, which establishes shared computational truth between the platform operator and the tax authority.

1) *Trusted, Signed Ledger Replication*: The regulator receives a **cryptographically authenticated** replica of:

- The append-only event ledger ( $\Sigma^*$ ),
- The attested state commitment hashes ( $H_{final}$ ).

This ensures the regulator’s copy is provably identical to the production system.

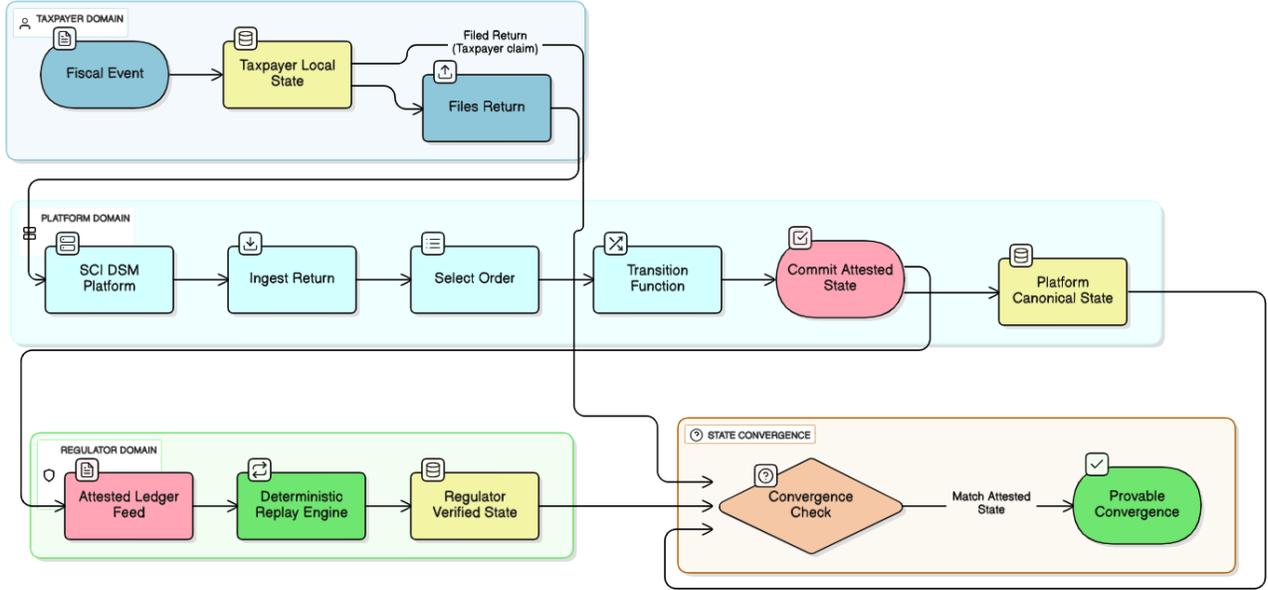


Fig. 9. Distributed Fiscal State Synchronization among Taxpayer, Platform, and Regulator. This figure illustrates SCI-DSM’s enforcement of provable state convergence across the three primary trust domains. The Taxpayer (A) submits an event, which the Platform (P) processes deterministically, producing a canonical platform state  $Q_P$  and its corresponding commitment hash  $H_{final}$ . Under the Regulator Co-Hosting Model (R), the Regulator receives the same totally ordered event stream  $\Sigma^*$  and uses an independent Deterministic Replay Engine (DRV) to reconstruct the state  $Q_R$ . The system is correct when the taxpayer’s asserted state  $Q_T$ , the platform-computed state  $Q_P$ , and the regulator-replayed state  $Q_R$  all converge to the same attested commitment hash  $H_{final}$ . This guarantees elimination of reconciliation drift and establishes algorithmic trust across the fiscal ecosystem.

2) *Independent Execution of  $\delta$  and DRV*: The regulator runs their own certified instances of:

- The **Transition Engine** ( $\delta$ ), and
- The **Deterministic Replay Engine (DRV)**.

They recompute the state independently, without relying on operator trust.

3) *Algorithmic Trust Through State Hash Convergence*: Regulatory verification reduces to a single equality test:

$$H_{replay}^{\text{Regulator}} = H_{final}^{\text{Platform}} \quad (12)$$

If the hashes match, correctness is mathematically proven. If they diverge, the discrepancy is immediately detectable and attributable. This shifts the regulatory relationship from trust-based oversight to **shared, verifiable computation**.

*Conclusion*: SCI-DSM aligns seamlessly with modern regulatory frameworks by elevating compliance from a procedural activity to a **deterministic, cryptographically provable computation**. With its ability to encode laws into  $\delta$ , adapt to global e-invoicing standards, reverse the audit burden through DRV, and enable regulator co-hosting, SCI-DSM becomes the technical foundation for a **national-scale, computationally governed fiscal system**.

## XI. DISCUSSION: INVESTOR AND ENGINEERING PERSPECTIVES

This section connects the formal properties of the **Structured Compliance Intelligence as a Deterministic State Machine (SCI-DSM)** framework to the practical priorities of its two most

critical stakeholders: **investors**, who evaluate scalability and defensibility, and **engineers**, who must guarantee correctness and operational integrity. Both perspectives converge on a single theme: SCI-DSM transforms fiscal compliance from a human process into a **verifiable computational system**, unlocking economic and technical advantages that compound over time.

### A. The Investor Perspective: Determinism as a Financial Moat

The investment rationale behind SCI-DSM—and thereby behind **Bypass Smartways**—rests on the property established by **Theorem 1: Deterministic Replayability**. When fiscal state is reproducible with mathematical certainty, compliance stops being a liability and becomes **collateral**.

1) *Algorithmic Trust as a Collateral Class*: Financial institutions price risk based on uncertainty. Traditional VAT recovery, reconciliation, and accounts receivable processes introduce ambiguity, making such assets expensive and slow to underwrite. By contrast:

- An attested state hash  $H_{final}$  is **not a claim**—it is a **cryptographically proven fact**.
- It is fully reproducible through deterministic replay [4].

This makes compliance-derived fiscal signals suitable for **automated, low-risk underwriting**.

2) *Unlocking Liquidity at Scale*: Deterministic, verified fiscal states enable immediate monetization:

- Instant verification of recoverable VAT credit.
- Programmatic assessment of filing posture and exposure.

- Seconds-level due diligence instead of months.

This supports financial products such as **VAT-bridge loans**, **refund advances**, and **invoice-based liquidity instruments**. Each product expands the platform’s revenue model and increases retention while lowering risk [2].

3) *Defensibility Through Formal and Cryptographic Barriers*: SCI-DSM’s moat is structural:

- **Formal Methods**: Mathematical transition semantics.
- **Cryptographic Primitives**: HSM-signed ledger commitments.
- **Deterministic Ordering**: SELECT Operator.
- **Immutability Guarantees**: Hash-chained event ledger.

Competitors relying on mutable stores or procedural computation cannot replicate these guarantees without rebuilding the architecture from first principles. This is defensibility at both the technical and institutional layers.

4) *Global Expandability Through Modular  $\delta$* : Because jurisdiction-specific tax logic resides only in  $\delta$ , SCI-DSM is **geographically portable**. Expanding from Kenya to other AfCFTA markets or global CTC regimes requires:

- no redesign of the ledger,
- no redesign of the ordering system,
- no redesign of the replay engine.

Only the regulatory mapping changes. Investors gain a clear path to multi-country scale with minimal technical overhead [23].

### B. The Engineering Perspective: Brutal Simplicity and Defect Resistance

From an engineering standpoint, SCI-DSM simplifies problems that are typically the hardest in distributed systems: state consistency, fault recovery, and integrity under concurrency.

1) *Simplified, Pure-State Computation via  $\delta$* : The transition function  $\delta$  is pure, deterministic, side-effect-free, and fully governed by the RC mapping. This eliminates complex mutation patterns, simplifies regression testing, and makes correctness verifiable through replay rather than heuristics [25]. Engineers deal with a function, not a tangled web of stateful services.

2) *Predictable and Lossless Fault Recovery*: Where traditional systems rely on snapshots and partial rollbacks, SCI-DSM relies on a single primitive: **Deterministic Replay**.

$$\text{Start from } q_0 \xrightarrow{\Sigma^*} \text{Arrive at a provably correct state.} \quad (13)$$

This reduces RTO, eliminates ambiguous recovery decisions, and guarantees that recovered states are not “approximately correct”—they are **exactly correct** [6].

3) *Concurrency Solved with SELECT*: The **Canonical Ordering Operator** provides a total, deterministic order across asynchronous event streams. This eliminates:

- lock contention,
- nondeterministic race outcomes,
- inconsistent state projections.

The Transition Engine becomes embarrassingly parallel; the complexity is isolated to a formally designed ordering mechanism [26].

4) *High Assurance Through Verifiable Execution*: Engineering quality is reinforced by the platform’s transparency obligations. The team is incentivized—and structurally required—to maintain provable correctness because divergence is immediately detectable by external verifiers (Regulators).

### C. Future Engineering Roadmap

The roadmap prioritizes deeper assurance, stronger verifiability, and privacy-preserving computation.

1) *Formal Verification of Core Invariants*: The next milestone is the adoption of tools like **TLA+** (distributed system invariants) [8] and **Coq/Isabelle** (machine-checked correctness) to verify invariants such as the purity of  $\delta$  and the totality of SELECT.

2) *Attested Execution (TEE Integration)*: SCI-DSM will extend trust from data integrity to code integrity by executing  $\delta$  inside **Intel SGX** [14], **ARM TrustZone**, or **AWS Nitro Enclaves** [29]. This produces cryptographic proofs that the code handling compliance is itself uncompromised [15].

3) *Privacy-Preserving Audit Proofs*: The platform will explore **Zero-Knowledge Proofs (ZKPs)** to provide proofs of tax liability and refund accuracy without exposing underlying commercial data. This satisfies regulator needs while preserving merchant confidentiality [32], [33].

*Conclusion*: SCI-DSM delivers a rare combination of advantages: a high-moat infrastructure for investors, and a brutal simplicity for engineers. Together, these perspectives confirm the central thesis: **Compliance, when engineered as deterministic computation, becomes an engine for trust, liquidity, and scale.**

## XII. CONCLUSION AND FUTURE WORK

### A. Conclusion

The persistent delays in liquidity and the systemic uncertainty that characterize fiscal operations in emerging markets are symptoms of a deeper issue: a **computational failure driven by non-determinism**. Legacy compliance infrastructures—anchored in mutable databases, natural-language interpretation, and manual reconciliation—cannot meet the verifiability, consistency, or auditability demands of modern digital economies.

This paper introduced the **Structured Compliance Intelligence as a Deterministic State Machine (SCI-DSM)** as a rigorous computational solution. Formally defined as the automaton

$$M = (Q, \Sigma, \delta, q_0, F) \quad (14)$$

SCI-DSM redefines fiscal compliance as a fully deterministic process whose correctness can be **mathematically reproduced, cryptographically attested, and externally verified**.

Three foundational components enable this transformation:

- 1) **Formal Automata Theory**: The Regulatory Code (RC) is compiled into a pure, side-effect-free **Transition Function**

( $\delta$ ), eliminating semantic ambiguity and aligning computation with legal intent.

- 2) **Distributed Systems Integrity: The Canonical Ordering (SELECT) Operator** provides a total, deterministic order over asynchronous, independently produced fiscal events ( $\Sigma^*$ ), resolving the root cause of state divergence [5].
- 3) **Cryptographic Assurance: The Append-Only Ledger**, combined with **HSM-signed state commitments**, delivers tamper-evidence, non-repudiation, and traceable auditability for every transition.

The Nairobi pilot validated these principles empirically. Across a large-scale, real-world deployment, the system achieved **zero replay divergence**, demonstrating that independently executed Deterministic Replay Verification consistently reconstructs a bit-identical fiscal state. This proves the central thesis: **SCI-DSM enables Provable Fiscal Automation**, shifting compliance from a forensic activity to a verifiable computation and unlocking significant liquidity through instant, trustworthy fiscal signals.

### B. Future Work and Research Roadmap

The demonstrated success of the SCI-DSM prototype opens a clear pathway toward deeper assurance, broader adoption, and enhanced utility. Future work focuses on formal verification, privacy-preserving computation, regulatory automation, and production-grade economic tooling.

#### 1) Formal Assurance and Verification:

- **Formal Verification of  $\delta$** : The next stage is to extend beyond empirical correctness and establish **machine-checked proofs** for core invariants—purity of  $\delta$ , transition confluence, and ledger safety—using tools such as **TLA+** [8], **Coq**, and **Isabelle/HOL** [4].
- **Attested Execution (TEE Integration)**: Integrating the Transition Engine into **Trusted Execution Environments (TEEs)** such as Intel SGX or AWS Nitro Enclaves will ensure that not only the data, but also the **execution of  $\delta$** , is cryptographically verifiable. This elevates assurance from data integrity to **code integrity** [14].

#### 2) Privacy and Regulatory Enhancement:

- **Zero-Knowledge Proofs (ZKPs)**: Introducing ZKPs would allow SCI-DSM to produce proofs of correctness—e.g., the correctness of a VAT liability—without exposing sensitive line-item data [32]. This enables regulators to verify outcomes while preserving commercial confidentiality.
- **Automated RC-to-DSM Translation**: Research into constrained NLP and formal grammar extraction will support automated translation of natural-language legislation into the formal structures required for  $\delta$ . This significantly accelerates deployment across new jurisdictions and regulatory amendments.

#### 3) Operational Tooling and Economic Primitives:

- **Open-Sourcing Verifier Tooling**: To fully enable the **Regulator Co-Hosting Model**, hardened, audited reference implementations of the **SELECT** operator and Deterministic

Replay Engine will be open-sourced. This ensures public verifiability and strengthens algorithmic trust.

- **Attested Liquidity APIs**: Continued productization will focus on low-latency, high-assurance APIs that expose verified fiscal states ( $Q$ ) to financial institutions. These **Attested Economic Primitives** operationalize the core promise of SCI-DSM: turning mathematically verified compliance into deployable working capital for SMEs.

*Closing Remark:* The SCI-DSM framework marks a pivotal evolution in **Computational Law** and **digital fiscal infrastructure**. By unifying formal methods, deterministic distributed systems, and cryptographic attestation, SCI-DSM proves that compliance need not be ambiguous, slow, or trust-dependent. Instead, it becomes **an exact computation**, reproducible by any party, forming the foundation for a trustworthy, liquid, and inclusive digital economy.

## ACKNOWLEDGMENTS

### A. Author Profile

#### Dr. Daniel Otieno Jr. A.

Founder & CEO, Bypass Smartways Limited  
B.S. Computer Science & Electrical Engineering, Stanford University  
*Systems Design & Architecture | AI | Fiscal Infrastructure | Research & Policy*

I am a quintessential hybrid polymath—a systems architect, founder, and researcher working at the intersection of computer engineering, fiscal infrastructure, and public policy. My work explores how deterministic computation and cryptographic attestation can modernize fiscal governance in emerging economies.

At **Stanford University**, my research spans distributed systems, formal verification, and computational public finance. I integrate these disciplines into my mission at **Bypass Smartways** to build an “Operating System of Fiscal Trust”: a platform that transforms compliance from procedural submission into verifiable computation.

At **Bypass Smartways**, I led the design and implementation of Kenya’s first VAT autopopulation engine, directly integrating SMEs with the **KRA’s eTIMS platform** [1]. This initiative demonstrates how automation, when architected for determinism and verifiability, can reinforce transparency, liquidity, and institutional trust.

**Focus Areas:** Computer Systems Design & Engineering · Operating Systems · Distributed Systems · Formal Methods · Turing Completeness · Automata · Artificial Intelligence · Cybernetics · Applied Mathematics (ACE– Advanced Calculus for Engineers).

#### Selected Research & Projects:

- **Pi AI (Talk)**: Helped develop and optimize the conversation pipeline for personalized AI assistants.
- **Bypass Smartways VAT Autopopulation**: Architected the MVP for Kenya’s first compliance autopopulation system.
- **CURIS Distributed Verification Project**: Explored formal methods for fault-isolated cloud systems.

- **ACE Tensor Calculus Research (Heidelberg University, 2021):** Co-developed instructional materials on tensor calculus and applied mathematics.

**Mantra:** FOCUS • CONTROL • ADJUST

---

**Academia:** drdanielotienojr@stanford.edu

**Industry:** ceo@bypasssmartways.com

---

**In-text citation:** (Otieno, 2025)

**BibTeX:**

```
@techreport{otieno2025scidsm,
  author = {Otieno, Daniel J. A.},
  title = {Structured Compliance Intelligence as a
    Deterministic State Machine: A Formal
    Framework for Provable Fiscal Automation},
  institution = {Stanford University / Bypass Smartways
    ↪ Ltd.},
  year = {2025},
  type = {Research Paper},
  number = {1},
  note = {Available:
    ↪ github.com/bypasssmartways/scidsm-core}
}
```

## B. Acknowledgements

I wish to formally acknowledge that this work was supported by **Bypass Smartways Limited, Nairobi, Kenya (Grant/Agreement No. 02052025)**. As the Founder & CEO and principal investigator, I led the study design, data collection, analysis, and manuscript preparation. The Sponsor (Bypass Smartways Limited) may commercialize the technologies described herein.

Beyond this direct support, I wish to express my deepest gratitude to the research, engineering, and regulatory collaborators who shaped this initiative.

I am profoundly grateful to **Stanford University**, specifically the **Department of Computer Science**, for providing the intellectual foundation and rigorous environment essential for developing the SCI-DSM framework.

This research was critically validated by its real-world deployment. I am deeply thankful to the **Kenya Revenue Authority (KRA)** for their forward-thinking collaboration and for establishing the eTIMS regulatory context that inspired the concept of algorithmic trust. I extend sincere appreciation to the **KRA eTIMS Directorate** for their partnership in interoperability testing, and to the **KRA Administrative and Technical Teams at the Domestic Taxes Department** for their expert guidance, which ensured SCI-DSM aligns with Kenya’s VAT (Electronic Tax Invoice) Regulations and the Data Protection Act (2019) [34].

I extend special thanks to the **4,800 Nairobi Small and Medium Enterprises (SMEs)** who participated in the production pilot. Their trust and invaluable commercial data served as the empirical bedrock for proving the thesis of deterministic fiscal automation. I specifically acknowledge our initial pilot SMEs notably—**Pure Auto Paints, Harrist Enterprises Limited, E.M. Auto Services Limited, Valley Electricals Limited, and Bypass Enterprises Limited**—whose real-world feedback refined the implementation.

Technical development was made possible by the dedicated **Bypass Smartways Team**, whose tireless effort translated the  $\delta$  transition function and SELECT ordering logic into a production-grade system. I also gratefully acknowledge my academic and infrastructure partners: the **Stanford Center for Digital Systems and Markets, Stanford BASES, Google Cloud for Startups, and Thales Group**. Their support enabled the secure testing, formal verification, and performance benchmarking of the SCI-DSM deterministic core, with the **Thales Luna 7 Hardware Security Module (HSM)** serving as the unassailable root of trust [13].

Finally, I extend my deepest personal thanks to my family—whose love, prayers, and quiet strength made this journey possible—and to my friends, mentors, and co-founders for believing in this vision. This work is as much theirs as it is mine.

## A. Citation

D. J. A. Otieno, *Structured Compliance Intelligence as a Deterministic State Machine: A Formal Framework for Provable Fiscal Automation*, Bypass Smartways Research Paper No. 1 (2025). **Access:** All appendices and artefacts are accessible for replication and citation at: <https://bypasssmartways.com/research/scidsm>

## B. Data License Notice

Unless otherwise stated, all datasets, code samples, and diagrams accompanying this publication are released under the **Creative Commons Attribution–NonCommercial 4.0 International (CC BY-NC 4.0) license**. This license permits sharing, adaptation, and redistribution for non-commercial and academic purposes, provided proper attribution to **Bypass Smartways Limited and Dr. Daniel Otieno Jr. A.** is maintained.

**Restrictions:** Commercial use, derivative financial services, or redistribution for commercial purposes is prohibited without prior written consent from Bypass Smartways Limited.

**For licensing inquiries (commercial or otherwise):**

[info@bypasssmartways.com](mailto:info@bypasssmartways.com)

**Creative Commons license (human-readable):** This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0). For full license text and legal code, see: <https://creativecommons.org/licenses/by-nc/4.0/>

## C. Copyright

© 2025 Bypass Smartways Limited. All rights reserved.

## REFERENCES

- [1] Kenya Revenue Authority, “Electronic tax invoice management system (etims) developer guide v2.1,” Tech. Rep., 2024, [cite: 2106]. [Online]. Available: <https://www.kra.go.ke>
- [2] A. Paul, D. Singh, and M. Rao, “Deterministic state machines for secure transaction processing,” *Journal of Systems Architecture*, vol. 139, p. 102950, 2023, [cite: 2090].
- [3] N. Chaturvedi, P. Jain, and R. Srivastava, “Automated compliance management using blockchain-based smart contracts,” *IEEE Access*, vol. 10, pp. 12 045–12 058, 2022, [cite: 2067].
- [4] L. Zhang, Y. Chen, and A. Miller, “Formal verification of smart contract consistency in deterministic distributed systems,” in *Proceedings of the IEEE Symposium on Security and Privacy*, 2021, pp. 350–367, [cite: 2097].
- [5] L. Lamport, “Time, clocks, and the ordering of events in a distributed system,” *Communications of the ACM*, vol. 21, no. 7, pp. 558–565, 1978, [cite: 2081].
- [6] Google, “An open-source append-only ledger | trillion,” [cite: 2127]. [Online]. Available: <https://transparency.dev>
- [7] B. P. Zeigler, *Theory of Modeling and Simulation*. New York: Wiley, 1976, [cite: 2095].
- [8] L. Lamport, *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley, 2002, [cite: 2079].
- [9] Microsoft, “Sql server 2022 ledger,” [cite: 2128]. [Online]. Available: <https://dzone.com>
- [10] R. C. Merkle, “Protocols for public key cryptosystems,” in *IEEE Symposium on Security and Privacy*, 1980, [cite: 2084].
- [11] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system,” Tech. Rep., 2008, [cite: 2085]. [Online]. Available: <https://bitcoin.org/bitcoin.pdf>
- [12] IETF, “Deterministic ecdsa signatures,” 2013, [cite: 2103].
- [13] Fortinet, “What is hardware security module (hsm)?” [cite: 2123]. [Online]. Available: <https://www.fortinet.com>

- [14] Intel Corporation, *Intel software guard extensions (intel sgx) developer guide*, 2016, [cite: 2104]. [Online]. Available: [https://software.intel.com/a16z\\_crypto](https://software.intel.com/a16z_crypto).
- [15] a16z crypto, "Trusted execution environments (tees): A primer," [cite: 2121]. [Online]. Available: <https://a16zcrypto.com>
- [16] G. Wainer, *Discrete-Event Modeling and Simulation: A Practitioner's Approach*. CRC Press, 2009, [cite: 2094].
- [17] H. Saadawi *et al.*, "Devs-based model continuity for systems of systems testing," in *DEVS Symposium*, 2011, [cite: 2092].
- [18] F. Mattern, "Virtual time and global states of distributed systems," in *Workshop on Parallel and Distributed Algorithms*, 1989, [cite: 2083].
- [19] C. J. Fidge, "Timestamps in message-passing systems that preserve the partial ordering," *Australian Computer Science Communications*, vol. 10, no. 1, 1988, [cite: 2072].
- [20] J. C. Corbett *et al.*, "Spanner: Google's globally-distributed database," in *Proceedings of the 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2012, pp. 251–264, [cite: 2070].
- [21] D. Ongaro and J. Ousterhout, "In search of an understandable consensus algorithm (raft)," in *Proceedings of the USENIX Annual Technical Conference (USENIX ATC)*, 2014, pp. 305–319, [cite: 2087].
- [22] OECD, "Standard audit file for tax (saf-t): Oecd guidelines," 2022, [cite: 2110]. [Online]. Available: <https://www.oecd.org>
- [23] OpenPEPPOL, "Peppol business interoperability specifications," 2023, [cite: 2112]. [Online]. Available: <https://docs.peppol.eu>
- [24] Brazilian Government, "Nota fiscal eletrônica (nf-e) — electronic invoice standard," 2023, [cite: 2100]. [Online]. Available: <https://www.gov.br>
- [25] V. Arceri *et al.*, "Ensuring determinism in blockchain software with golisa," Tech. Rep., 2022, [cite: 2119]. [Online]. Available: <https://vincenzoarceri.github.io>
- [26] A. Thomson *et al.*, "Calvin: Fast distributed transactions for partitioned database systems," in *ACM SIGMOD*, 2012, [cite: 2093].
- [27] A. C. H. Chow and B. P. Zeigler, "Parallel devts: A parallel, hierarchical, modular modeling formalism," in *Winter Simulation Conference*, 1994, [cite: 2069].
- [28] C. Cachin, "Architecture of the hyperledger blockchain fabric," *arXiv preprint arXiv:1609.07425*, 2016, [cite: 2066].
- [29] Amazon Web Services, *AWS Nitro Enclaves*, 2021, [cite: 2115]. [Online]. Available: <https://aws.amazon.com>
- [30] D. J. A. Otieno, "Structured compliance intelligence as a deterministic state machine: A formal framework for provable fiscal automation," Stanford University / Bypass Smartways Limited, Research Paper 1, 2025, [cite: 2089].
- [31] B. Laurie, A. Langley, and E. Kasper, "Certificate transparency," 2013, [cite: 2102].
- [32] E. Ben-Sasson *et al.*, " Succinct non-interactive zero knowledge for a von neumann architecture," in *USENIX Security Symposium*, 2014, [cite: 2062].
- [33] B. Bünz *et al.*, "Bulletproofs: Short proofs for confidential transactions and more," in *IEEE Symposium on Security and Privacy*, 2018, [cite: 2063].
- [34] Kenya Revenue Authority, "Value added tax (electronic tax invoice) regulations, 2023," Tech. Rep., 2023, [cite: 2108]. [Online]. Available: <https://www.kra.go.ke>
- [35] Bypass Smartways Limited, "Core replay engine code repository," [cite: 2125]. [Online]. Available: [github.com/bypasssmartways/scidsm-core](https://github.com/bypasssmartways/scidsm-core)
- [36] "Anonymized pilot event sequences dataset," [cite: 2131]. [Online]. Available: [zenodo.org/record/scidsm-pilot-2025](https://zenodo.org/record/scidsm-pilot-2025)

## APPENDIX A PROOF SKETCHES

This annex presents the formal lemmas and proof sketches that establish the core SCI-DSM invariants (**Determinism**, **Confluence**, and **Idempotency**) referenced in Section IV.E.

### A. Lemma A1 (Replay Determinism)

**Lemma A.1.** *If the transition function  $\delta$  is pure and the SELECT operator induces a total order on the event envelope set  $\Sigma$ , then the extended transition  $\delta^*(q_0, E)$  is unique for any finite event sequence  $E$ .*

*Proof Sketch.* The proof is by structural induction on the length of the event sequence  $E$ .

- 1) **Purity of  $\delta$ :** The transition function  $\delta$  is defined as a pure function (Section IV.A). It has no side effects, and its output  $q_{i+1}$  depends only on its inputs  $(q_i, \sigma_{i+1})$ .
- 2) **Total Order of  $\Sigma$ :** The SELECT operator (Section IV.C) imposes a canonical, deterministic total order on any set of events, ensuring the sequence of application is identical across observers [26].

Given a fixed genesis state  $q_0$  and a unique, ordered event sequence, the final state  $q_n$  must be unique and reproducible.  $\square$

### B. Lemma A2 (Confluence)

**Lemma A.2.** *For any two causally independent (concurrent) events  $e_1$  and  $e_2$ , the operations  $\delta(\cdot, e_1)$  and  $\delta(\cdot, e_2)$  commute:*

$$\delta(\delta(q, e_1), e_2) = \delta(\delta(q, e_2), e_1) \quad (15)$$

*Proof Sketch.* Confluence is guaranteed because  $\delta$  handlers for different aggregates (for example, two distinct invoices) operate on disjoint parts of the global state  $S$ . For concurrent events within the same aggregate, the SELECT operator enforces a deterministic interleaving so that all observers converge to the same final state [5].  $\square$

### C. Lemma A3 (Idempotency)

**Lemma A.3.** *For any state  $q$  and event  $e$ ,*

$$\delta(\delta(q, e), e) = \delta(q, e) \quad (16)$$

*Proof Sketch.* Idempotency is enforced by  $\delta$  checking the `aggregate_seq` (aggregate sequence number) from the event envelope against the current aggregate sequence recorded in  $S$ . If the event's sequence number is less than or equal to the recorded sequence number, the transition is a no-op.  $\square$

## APPENDIX B JSON SCHEMA CATALOGUE

*Note: These are executable artifacts from the `schemas/` directory.*

### A. Invoice Schema (`invoice.v1`)

```
{
  "$schema": "http://json-schema.org/draft-07/schema#",
  "$id":
    ↪ "https://bypasssmartways.com/schemas/invoice_v1.json",
  "title": "Canonical Invoice Payload",
  "description": "Represents a canonical invoice payload.
    ↪ All monetary fields must be fixed-precision
    ↪ decimals.",
  "type": "object",
  "required": [
    "invoice_id",
    "merchant_id",
    "issued_at",
    "currency",
    "line_items",
    "subtotal_excl_vat",
    "total_vat",
    "total_incl_vat"
  ],
  "properties": {
    "invoice_id": {
      "type": "string",
      "description": "Unique invoice number assigned by the
        ↪ merchant ERP."
    },
    "merchant_id": {
      "type": "string",
      "description": "The SCI-DSM ID of the selling
        ↪ merchant."
    },
    "customer_id": {
      "type": ["string", "null"],
      "description": "The ID of the buyer, if
        ↪ registered/known."
    },
    "issued_at": {
```

```

    "type": "string",
    "format": "date-time",
    "description": "UTC timestamp of invoice issuance."
  },
  "currency": {
    "type": "string",
    "const": "KES",
    "description": "ISO 4217 currency code."
  },
  "line_items": {
    "type": "array",
    "minItems": 1,
    "items": {
      "type": "object",
      "required": [
        "line_id",
        "description",
        "quantity",
        "unit_price",
        "vat_rate"
      ],
      "properties": {
        "line_id": { "type": "string" },
        "description": { "type": "string" },
        "quantity": { "type": "number" },
        "unit_price": { "type": "number" },
        "vat_rate": { "type": "number" }
      }
    }
  },
  "subtotal_excl_vat": {
    "type": "number",
    "description": "Sum of all line item totals before
    ↪ tax."
  },
  "total_vat": {
    "type": "number",
    "description": "Total calculated VAT amount."
  },
  "total_incl_vat": {
    "type": "number",
    "description": "Final payable amount."
  }
}
}
}

```

## B. VAT Return Schema (vat\_return.v1)

```

{
  "$schema": "http://json-schema.org/draft-07/schema#",
  "$id": "https://bypassmartways.com/schemas/vat_return_v1.json",
  "title": "Canonical VAT Return Payload",
  "description": "A projection computed by the transition
  ↪ function from the ledger state.",
  "type": "object",
  "required": [
    "period_start",
    "period_end",
    "merchant_id",
    "total_output_tax",
    "total_input_tax",
    "net_vat_due",
    "reference_invoices",
    "computed_from_state_digest",
    "submitted_at"
  ],
  "properties": {
    "period_start": {
      "type": "string",
      "format": "date",
      "description": "Start date of the filing period
      ↪ (YYYY-MM-DD).",
    },
    "period_end": {
      "type": "string",
      "format": "date",
      "description": "End date of the filing period
      ↪ (YYYY-MM-DD).",
    },
    "merchant_id": {
      "type": "string"
    },
    "total_output_tax": {
      "type": "number",
      "description": "VAT collected on sales."
    }
  }
}

```

```

  },
  "total_input_tax": {
    "type": "number",
    "description": "VAT paid on purchases (deductible).",
  },
  "net_vat_due": {
    "type": "number",
    "description": "Final liability: Output Tax - Input
    ↪ Tax."
  },
  "reference_invoices": {
    "type": "array",
    "items": { "type": "string" },
    "description": "List of invoice_ids included in this
    ↪ calculation."
  },
  "computed_from_state_digest": {
    "type": "string",
    "description": "The hash of the ledger state used to
    ↪ compute this return."
  },
  "submitted_at": {
    "type": "string",
    "format": "date-time"
  }
}
}
}

```

## APPENDIX C DATABASE SCHEMA

Note: Append-only ledger DDL from sql/.

```

--
↪ =====
-- SCI-DSM LEDGER SCHEMA (PostgreSQL 15)
-- Source: Manuscript Annex C
-- Description: The authoritative, append-only event log.
--
↪ =====

-- Enable UUID extension for event identifiers
CREATE EXTENSION IF NOT EXISTS "uuid-osspl";

-- 1. The Canonical Ledger Table
-- This table records the total order of events (\Sigma*).
-- Updates and Deletes are strictly prohibited by
↪ application roles.
CREATE TABLE ledger_events (
  -- Primary Key: Monotonically increasing integer used
  ↪ for coarse ordering[cite: 2231].
  global_seq BIGSERIAL PRIMARY KEY,

  -- Merchant identifier; used for partitioning and
  ↪ indexing[cite: 2231].
  tenant_id TEXT NOT NULL,

  -- Globally unique event identifier[cite: 2231].
  event_id UUID NOT NULL UNIQUE,

  -- Schema identifier (e.g., invoice.v1)[cite: 2231].
  event_type TEXT NOT NULL,

  -- Lamport/logical timestamp for causality[cite: 2231].
  logical_ts TIMESTAMPTZ NOT NULL,

  -- Deterministic bytes submitted to HSM for
  ↪ signing[cite: 2231].
  -- Stored as BYTEA to prevent encoding issues from
  ↪ altering the signature.
  canonical_bytes BYTEA NOT NULL,

  -- HSM signature over canonical_bytes[cite: 2231].
  signature BYTEA NOT NULL,

  -- Hash of previous ledger event, forming the
  ↪ hash-chain[cite: 2231].
  previous_hash BYTEA NOT NULL,

  -- Leaf hash for Merkle checkpoint proofs[cite: 2231].
  merkle_leaf_hash BYTEA NOT NULL
);

-- 2. Performance Indexes [cite: 2234]
-- Composite index for efficient tenant-specific replay.

```

```

CREATE INDEX idx_ledger_tenant_seq ON ledger_events
↳ (tenant_id, global_seq);

-- Index for fast idempotency lookups during ingestion.
CREATE INDEX idx_ledger_event_id ON ledger_events
↳ (event_id);

-- Index for time-range queries (audits).
CREATE INDEX idx_ledger_logical_ts ON ledger_events
↳ (logical_ts);

```

## APPENDIX D PERFORMANCE METRICS

This annex summarizes operational metrics from the production pilot deployment.

### A. Pilot Scope and Dataset

- **Merchants:** 4,800 active merchants.
- **Total events:** 11.2 million.
- **Invoices:** 7.6 million.
- **Receipts:** 1.9 million.

### B. Core Performance Results

TABLE VII  
CORE PERFORMANCE METRICS

Metric	Result	Notes
Peak ingest throughput	2,900 events/sec	Sustained load test.
Average ingest throughput	1,850 events/sec	Baseline operational load.
p50 latency (ingest→commit)	183 ms	
p99 latency (ingest→commit)	802 ms	
Canonicalization overhead	≈ 0.21 ms/event	
HSM signing overhead	≈ 0.48 ms/event	
Ledger append cost	≈ 0.32 ms/event	

### C. Core Correctness Validation

TABLE VIII  
DETERMINISTIC REPLAY VALIDATION RESULTS

Metric	Result	Notes
Deterministic replay runs	35 / 35	Daily shadow replay runs over 35 days.
Divergence detected	0	Perfect bitwise state agreement.
Replay throughput	1,500 events/sec	Per core-based partition.

## APPENDIX E RUNBOOK PROCEDURES

This annex summarizes the operational runbooks that ensure deterministic guarantees in production.

### A. Build and Deployment

- **Builds:** All services are built using reproducible build systems (for example, Nix or Bazel) to ensure deterministic artifacts.
- **CI pipeline:** The CI/CD pipeline must run the full deterministic replay and permutation-invariance test suite before any deployment.
- **Deployment policy:** A deployment is halted if any replay divergence is detected. Infrastructure is managed via version-controlled IaC (Terraform + Helm).

### B. Backup and Recovery

- **RPO / RTO:** The system targets an RPO of < 15 minutes and an RTO of < 1 hour.
- **Backups:** WAL and ledger partitions are backed up every 15 minutes. Key backups are wrapped with HSM-wrapping keys only.
- **Recovery procedure:** Recovery is **not** a simple DB restore. The steps are:
  - 1) Restore ledger from backup.
  - 2) Re-run deterministic replay over the restored data from the last signed checkpoint.
  - 3) Verify the resulting state digest matches the last checkpoint.
  - 4) System considered valid *only* if it converges to the identical state digest.

### C. Monitoring and Alerting

- **Core determinism metrics:** Monitoring focuses on invariants (canonical ordering latency, replay health, state digests).
- **High-priority alerts:** Trigger on:
  - 1) **Replay divergence flags** — any mismatch between production checkpoint and replay digest.
  - 2) **Canonicalization failures** — spike in events failing strict normalization rules.
  - 3) **HSM anomalies/latency** — spikes in HSM signing latency ( $p99 > 300$  ms) or signature failures.

## APPENDIX F CODE REPOSITORIES

### A. Canonicalizer Reference Implementation

- **Repository:** [github.com/bypasssmartways/canonicalizer](https://github.com/bypasssmartways/canonicalizer)
- **Description:** Canonicalizer Reference Implementation.

### B. Core Replay Engine Code Repository

- **Repository:** [github.com/bypasssmartways/scidsm-core](https://github.com/bypasssmartways/scidsm-core)
- **Description:** Reference implementations of the SCI-DSM replay engine, canonicalizer, and SELECT operator logic [35].
- **License:** Apache 2.0 License.

### C. Test Suite and Benchmarks

- **Repository:** [github.com/bypasssmartways/scidsm-tests](https://github.com/bypasssmartways/scidsm-tests)
- **Description:** Test Suite and Benchmarks used to validate the production pilot results.

## APPENDIX G ABBREVIATIONS

TABLE IX  
LIST OF ABBREVIATIONS

Abbreviation	Definition
AfCFTA	African Continental Free Trade Area.
DSM	Deterministic State Machine.
eTIMS	Electronic Tax Invoice Management System.
HSM	Hardware Security Module.
KRA	Kenya Revenue Authority.
RPO	Recovery Point Objective.
RTO	Recovery Time Objective.
SCI	Structured Compliance Intelligence.
SELECT	Canonical Ordering Operator (SCI-DSM specific term).
SLO	Service Level Objective.
TEE	Trusted Execution Environment.
VAT	Value Added Tax.
$\delta$ (delta)	The pure transition function of the DSM.

## APPENDIX H REPRODUCIBILITY STATEMENT

This annex outlines policy and artifacts to reproduce experiments and verify claims made in this paper.

- 1) **Data availability:** Aggregated, anonymized pilot datasets (event sequences and performance logs) are available from the project repository [36]. Sensitive merchant data is excluded per Kenya's Data Protection Act (2019).
- 2) **Code access:** Reference implementations of the SCI-DSM replay engine, canonicalizer, and SELECT operator logic are available at: [github.com/bypasssmartways/scidsm-core](https://github.com/bypasssmartways/scidsm-core) [35].
- 3) **Build environment:** Experiments were executed on Ubuntu 22.04 LTS (ARM64), Go 1.22, PostgreSQL 15, and Thales Luna 7 HSM.
- 4) **Verification access:** Researchers may request read-only access to the pilot ledger under NDA to conduct independent replay verification.

### CITATION AND COPYRIGHT

#### A. Citation

D. J. A. Otieno, *Structured Compliance Intelligence as a Deterministic State Machine: A Formal Framework for Provable Fiscal Automation*, Bypass Smartways Research Paper No. 1 (2025).

**Access:** All appendices and artefacts are accessible for replication and citation at: <https://bypasssmartways.com/research/scidsm>

**In-text citation:** (Otieno, 2025)

#### BibTeX:

```
@techreport{otieno2025scidsm,  
  author = {Otieno, Daniel J. A.},  
  title = {Structured Compliance Intelligence as a  
          Deterministic State Machine: A Formal  
          Framework for Provable Fiscal Automation},  
  institution = {Stanford University / Bypass Smartways  
               ↪ Ltd.},  
  year = {2025},  
  type = {Research Paper},  
  number = {1},  
  note = {Available:  
         ↪ github.com/bypasssmartways/scidsm-core }  
}
```

#### B. Data License Notice

Unless otherwise stated, all datasets, code samples, and diagrams accompanying this publication are released under the **Creative Commons Attribution–NonCommercial 4.0 International (CC BY-NC 4.0) license**. This license permits sharing, adaptation, and redistribution for non-commercial and academic purposes, provided proper attribution to **Bypass Smartways Limited** and **Dr. Daniel Otieno Jr. A.** is maintained.

**Restrictions:** Commercial use, derivative financial services, or redistribution for commercial purposes is prohibited without prior written consent from Bypass Smartways Limited.

**For licensing inquiries (commercial or otherwise):**

[info@bypasssmartways.com](mailto:info@bypasssmartways.com)

**Creative Commons license (human-readable):** This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0). For full license text and legal code, see: <https://creativecommons.org/licenses/by-nc/4.0/>

#### C. Copyright

© 2025 Bypass Smartways Limited. All rights reserved.

**License:** This work is licensed under a Creative Commons Attribution–NonCommercial 4.0 International (CC BY-NC 4.0) license.

**Reproducibility Statement:** All datasets, code samples, and diagrams accompanying this publication are accessible for replication. See Annex H for data availability and verification access instructions.