
RESEARCH ARTICLE

Tax digital twin: A reference architecture for continuous enterprise tax compliance

Manish Kumar¹ and Manykandaprebou Vaitinadin²

¹ *Independent Researcher, NC, USA*

² *Independent Researcher, CA, USA*

Corresponding Author: Manish Kumar, **E-mail:** rajpmmanish@gmail.com

ABSTRACT

Enterprise tax compliance remains predominantly reactive despite significant advances in enterprise information systems. Existing tax technologies primarily identify errors after transactions have been executed, limiting organizations' ability to anticipate how strategic business decisions—such as market expansion, supply-chain restructuring, or regulatory change—will affect future tax obligations. This paper addresses this limitation by developing a Tax Digital Twin (TDT) reference architecture that enables continuous simulation of enterprise tax compliance before operational decisions are implemented. Following a Design Science Research (DSR) methodology, the study introduces a reusable architectural artifact that extends digital twin theory beyond its traditional application to physical assets into the domain of enterprise tax management. Building upon Grieves' physical–virtual twin paradigm and control-theoretic synchronization principles, the proposed architecture represents enterprise tax posture as a continuously synchronized virtual model capable of predictive analysis, scenario simulation, divergence detection, and decision support. The architecture is formalized through a state representation model, synchronization mechanisms, and divergence-driven governance processes, and is illustrated using four representative enterprise scenarios involving multi-jurisdiction expansion, exemption certificate expiration, tax engine migration, and post-merger integration. Rather than replacing existing AI-based tax detection systems, the proposed Tax Digital Twin complements them by introducing proactive compliance simulation and enterprise decision intelligence. Based on a structured review of the digital twin and enterprise tax compliance literatures, no prior study was identified that proposes a reference architecture for continuous enterprise tax compliance simulation using digital twin principles; this paper addresses that gap and establishes a theoretical foundation for future empirical research, AI-enabled tax systems, and next-generation enterprise information systems.

KEYWORDS

Digital Twin; Enterprise Information Systems; Enterprise Tax Compliance; Design Science Research; Enterprise Architecture; ERP Systems; Predictive Compliance

ARTICLE INFORMATION

ACCEPTED: 20 October 2025

PUBLISHED: 09 November 2025

DOI: 10.32996/jefas.2025.7.7.3

1 Introduction

Enterprise organizations increasingly rely on Enterprise Resource Planning (ERP) systems to automate financial, operational, and regulatory processes. Despite significant advances in enterprise information systems, enterprise tax compliance remains predominantly reactive. Tax obligations are typically evaluated only after business transactions have been executed: an order is processed, tax determination rules are applied, accounting entries are generated, and only subsequently do reconciliation procedures, audits. While these technologies have significantly improved post-transaction tax monitoring, they share a common structural limitation—they evaluate transactions that have already occurred rather than predicting the consequences of decisions that have not yet been made.

This reactive paradigm limits the ability of organizations to answer strategic, forward-looking questions that increasingly influence enterprise decision-making. For example, how will opening a new distribution center affect economic nexus obligations

over the next six months? What tax exposure will result from restructuring an international supply chain? Which exemption certificates are likely to expire before the next renewal cycle, and what financial risk will that create? Existing enterprise tax technologies are designed primarily for transaction validation and compliance reporting; they provide limited support for continuously simulating future tax posture before operational or strategic decisions are implemented.

The ability to simulate future enterprise states is precisely the problem that digital twin technology was originally developed to address. A digital twin is a continuously synchronized virtual representation of a real-world system that supports monitoring, prediction, optimization, and scenario simulation throughout its lifecycle (Grieves, 2014; Grieves & Vickers, 2017). Originally introduced for product lifecycle management and manufacturing systems, digital twins have since been successfully applied to smart factories, industrial operations, transportation, healthcare, and supply-chain resilience (Tao et al., 2018; Ivanov & Dolgui, 2021). Their common characteristic is the maintenance of a continuously updated virtual model capable of predicting the consequences of future actions before they occur in the physical environment.

Despite rapid advances in digital twin research, their application within enterprise tax compliance remains largely unexplored. Existing AI-driven tax systems—including recent architectures for AI-assisted tax anomaly detection (Kumar, 2026)—focus primarily on identifying tax errors, compliance violations, or suspicious transactions after business events have been executed. Comparatively little research has investigated how enterprise information systems can maintain a continuously synchronized virtual representation of an organization's tax posture that supports predictive compliance analysis, scenario simulation, and proactive enterprise decision-making.

This paper addresses this research gap by proposing a Tax Digital Twin (TDT), a reference architecture for continuous enterprise tax compliance simulation. Following a Design Science Research (DSR) methodology (Hevner et al., 2004; Peffers et al., 2007), the proposed Tax Digital Twin constitutes a reusable architectural artifact that extends digital twin theory beyond physical assets into the domain of enterprise regulatory compliance. Building upon Grieves' physical–virtual twin paradigm and control-theoretic synchronization principles, the architecture models enterprise tax posture as a continuously evolving virtual entity capable of predictive simulation, divergence detection, and decision support.

The proposed architecture incorporates the unique characteristics of enterprise tax processes, including jurisdiction-dependent taxability rules, continuously evolving nexus obligations, time-sensitive exemption certificates, changing regulatory requirements, and heterogeneous enterprise data originating from ERP, procurement, finance, contract management, human resources, and facilities management systems. Unlike traditional tax engines that evaluate transactions individually, the Tax Digital Twin continuously models the overall compliance state of the enterprise and evaluates how future business decisions influence regulatory exposure before operational execution.

This study makes four principal contributions to the Enterprise Information Systems literature.

- C1. Based on a structured literature search (Section 2.1), it introduces, to the authors' knowledge, the first Tax Digital Twin reference architecture for representing enterprise tax compliance as a continuously synchronized virtual enterprise model.
- C2. It extends digital twin theory and control-theoretic synchronization principles from physical systems to enterprise regulatory processes by formalizing tax state representation, synchronization dynamics, and divergence detection.
- C3. It develops a reusable enterprise architecture comprising layered data integration, simulation services, governance mechanisms, and predictive compliance capabilities that support proactive enterprise decision-making.
- C4. It demonstrates the applicability of the proposed architecture through representative enterprise scenarios—including multi-jurisdiction expansion, exemption certificate lifecycle management, tax engine migration, and post-merger integration—and derives design knowledge intended to guide future research on AI-enabled enterprise information systems.

Beyond enterprise taxation, the proposed architecture illustrates how digital twin concepts can support continuous regulatory intelligence within enterprise information systems. The underlying architectural principles—including continuous synchronization, predictive simulation, divergence-driven governance, and AI-enabled decision support—are applicable to broader enterprise domains such as procurement compliance, financial governance, sustainability reporting, and regulatory risk management.

The remainder of this paper is organized as follows. Section 2 reviews the literature on digital twins, enterprise information systems, and AI-enabled tax compliance. Section 3 develops the theoretical foundations and Design Science perspective underpinning the proposed artifact. Section 4 presents the Tax Digital Twin reference architecture together with its formal state

representation and synchronization model. Section 5 illustrates the architecture through representative enterprise scenarios. Section 6 discusses theoretical implications, implementation considerations, limitations, and future research directions. Section 7 concludes the paper.

2. Literature Review

2.1 Literature Review Methodology

To substantiate the novelty claims advanced in this paper, a structured search was conducted across academic and preprint sources (Google Scholar-indexed literature, ScienceDirect, arXiv, and Research Square) covering publications through mid-2026, combining three concept groups: (i) "digital twin" OR "digital shadow" OR "cyber-physical system"; (ii) "tax compliance" OR "tax administration" OR "enterprise tax" OR "regulatory compliance"; and (iii) "enterprise information system" OR "ERP" OR "enterprise architecture." The search identified related but structurally distinct prior work. Xu, Bai, Zhou, and Huang (2025) describe government-side tax administration systems in China's Golden Tax Project III that "map the business objects of tax administration as digital twins," but do not propose a reference architecture and address the tax authority's perspective rather than the taxpayer enterprise's. Gu, Tian, Yang, and Liu (2025) propose Reg-Twin, a digital twin architecture for financial regulatory compliance monitoring incorporating a policy rule engine and drift monitoring, representing the closest structural precedent identified, though applied to financial and anti-money-laundering compliance rather than tax. Zhang, Bahsoon, Tziritas, and Theodoropoulos (2024) establish a similar physical-virtual-decision architecture for adaptive compliance in cyber-physical systems, applied to industrial and warehouse logistics contexts. No identified work proposes a digital twin reference architecture specifically for enterprise-side, ERP-integrated tax compliance simulation, which is the gap this paper addresses.

2.2 Origins and Evolution of Digital Twin Theory

The digital twin concept originated in manufacturing and aerospace contexts. (Grieves, 2014) introduced the term to describe a three-part model consisting of a physical product, a virtual product, and the connections binding the two, developed initially within product lifecycle management. (Glaessgen and Stargel, 2012), working independently within NASA and U.S. Air Force vehicle contexts, articulated a closely related paradigm centered on ultra-high-fidelity simulation for predicting the behavior of physical assets throughout their operational life. (Grieves and Vickers, 2017) later extended the concept explicitly to complex systems, framing the digital twin as a mechanism for anticipating and mitigating unpredictable, undesirable emergent behavior — a framing directly relevant to tax compliance, where non-compliance frequently emerges from the interaction of otherwise individually reasonable business decisions rather than from any single erroneous transaction.

Subsequent literature has converged on a small number of structural models for what a digital twin comprises. (Tao, Zhang, Liu, and Nee, 2018) proposed a five-dimension model — physical entity, virtual entity, twin data, services, and connections — that has become a standard reference structure in the manufacturing digital twin literature and is adopted, with domain-specific reinterpretation, as the organizing structure for the architecture proposed in Section 4. (Kritzinger, Karner, Traar, Henjes, and Sihn, 2018) contributed an influential distinction between digital models, digital shadows, and true digital twins, based on the directionality and automation of data flow between physical and virtual entities — a distinction this paper relies on in Section 3 to clarify what differentiates a genuine tax twin from a monitoring dashboard. Systematic reviews by (Jones, Snider, Nassehi, Yon, and Hicks, 2020), (Fuller, Fan, Day, and Barlow, 2020), Qi, Tao, Zuo, and Zhao (2019), and Barricelli, Casiraghi, and Fogli (2019) collectively document the concept's rapid diffusion across manufacturing, healthcare, and infrastructure domains, while also noting that enabling technologies and evaluation methodology remain considerably more mature than domain-specific theoretical grounding outside of engineering contexts.

2.3 Digital Twins Beyond Manufacturing: Supply Chain and Enterprise Systems

The clearest precedent for extending digital twin theory into enterprise, rather than purely physical-engineering, contexts comes from supply chain management. Ivanov and Dolgui (2021) proposed the digital supply chain twin, a computerized model representing the state of an entire supply network at any given moment, capable of simulating the propagation of disruption risk and supporting resilience decisions before disruptions materialize. Ivanov, Dolgui, Das, and Sokolov (2019) extended this framing to data-driven optimization and visibility across the ripple effects of supply chain disruption. Notably for the present paper's theoretical grounding, Ivanov, Sethi, Dolgui, and Sokolov (2018) explicitly connect digital twin and supply chain simulation research to control theory, arguing that operational systems facing continuous disturbance are naturally modeled as closed-loop control systems with reference states, feedback, and corrective action — a framing this paper adopts and extends into the tax domain in Section 3.

This supply-chain literature establishes an important precedent: that digital twin theory can be meaningfully extended from physical, sensor-instrumented systems to enterprise-level systems whose “state” is substantially informational and financial rather than purely physical. Enterprise tax posture — nexus status, exemption certificate validity, applicable rates, filing obligations — is precisely this kind of informational-financial state: it is grounded in physical and transactional reality but is itself an abstraction computed from underlying data rather than a directly sensed physical quantity.

Table 1. Comparison of Digital Twin applications across domains and the proposed Tax Digital Twin architecture.

Study	Application Domain	Physical Entity	Virtual Representation	Continuous Synchronization	Predictive Simulation	Enterprise Tax Compliance
Grievies & Vickers (2017)	Manufacturing	✓	✓	✓	✓	X
Tao et al. (2018)	Smart Manufacturing	✓	✓	✓	✓	X
Ivanov & Dolgui (2021)	Supply Chain	Partial	✓	✓	✓	X
This Paper (TDT)	Enterprise Tax	✓	✓	✓	✓	✓

2.4 AI-Driven Tax Compliance Systems and the Detection–Simulation Gap

A growing body of applied work addresses AI-driven detection of tax non-compliance within enterprise and e-commerce contexts. Institutional guidance from the OECD (2020) on digital transformation in tax administration, and the well-documented expansion of economic nexus obligations following the Supreme Court’s ruling in *South Dakota v. Wayfair, Inc.* (2018), both reflect and reinforce this detection-and-remediation posture as the dominant paradigm in both practitioner and academic treatments of AI in tax compliance.

What is absent from this literature is a system capable of answering a counterfactual or prospective question about tax posture — not “was this transaction handled correctly” but “what will our tax posture look like under a hypothetical future state.” This is not a shortcoming of the detection literature, which addresses a genuinely different and equally necessary problem; rather, it identifies an adjacent, largely unaddressed problem for which digital twin theory, as developed in the supply chain and manufacturing literature reviewed above, offers an appropriate and largely untested theoretical foundation.

Unlike manufacturing digital twins, enterprise regulatory systems do not synchronize physical measurements obtained from sensors. Instead, they synchronize semantic enterprise state derived from transactional systems, business rules, organizational structures, and regulatory knowledge. Consequently, extending digital twin theory into enterprise tax compliance requires new abstractions for representing regulatory state, synchronization semantics, temporal obligations, and governance processes.

2.5 Research Gap

Synthesizing the literatures reviewed above: digital twin theory has been substantively extended from physical manufacturing systems into enterprise supply chain contexts, demonstrating that the underlying physical–virtual synchronization model generalizes beyond sensor-instrumented physical assets. Separately, AI-driven tax compliance research has matured considerably in detection and risk-scoring capability but remains structurally oriented toward retrospective transaction evaluation. No prior work, to the author’s knowledge, has proposed a digital twin architecture specifically for enterprise tax processes, nor formalized the state representation, synchronization dynamics, and simulation services such a twin would require. This paper addresses that gap.

Table 2. Research gaps identified in the literature and corresponding contributions of the proposed Tax Digital Twin.

Existing Research Stream	Current Limitation	Contribution of this Paper
Digital Twin Research	Focused primarily on physical assets	Extends Digital Twin theory to enterprise tax compliance
Enterprise Tax Systems	Reactive transaction processing	Continuous enterprise tax-state representation

Existing Research Stream	Current Limitation	Contribution of this Paper
AI Tax Detection	Detects historical anomalies	Predictive tax-state simulation
ERP Tax Engines	Transaction-level tax calculation	Enterprise-wide tax posture management
Enterprise Architecture	Limited support for predictive compliance	Six-layer Tax Digital Twin reference architecture

3. Theoretical Foundations

3.1 The Physical–Virtual Twin Model Applied to Tax

Following (Grieves and Vickers, 2017), a digital twin comprises three elements: a physical entity, a virtual entity, and the data connections linking them. Mapped onto enterprise tax processes, the physical entity is the enterprise's actual operational and transactional reality — its ERP transaction stream, its legal entity structure, its contracts, and the physical operations (warehouses, personnel, inventory) that give rise to nexus and taxability determinations. The virtual entity is a formal, computable model of the enterprise's tax posture — not a record of what happened, but a live representation of what the enterprise's tax obligations currently are and, critically, what they would become under a hypothetical future state. The twin data connection is the continuously updated dataset and synchronization protocol binding the two: the mechanism by which the virtual entity's model is kept faithful to the physical entity's actual state.

Kritzinger, Karner, Traar, Henjes, and Sihn's (2018) distinction between digital models, digital shadows, and digital twins is directly relevant here. A digital model of tax posture — a static compliance report, for instance — involves no automated data flow in either direction. A digital shadow involves automated flow from the physical entity to the virtual entity (the virtual model updates automatically as transactions occur) but no return flow: this describes most existing tax dashboards and even most detection systems, which observe but do not feed conclusions back into a governing model. A true digital twin requires bidirectional automated flow: the virtual entity updates from physical data, and outputs of the virtual entity — simulated projections, divergence alerts — are capable of informing physical-entity decisions and, in principle, of triggering automated corrective action in the physical system. The architecture proposed in Section 4 is designed to satisfy this stronger bidirectional criterion, distinguishing it from monitoring dashboards that are sometimes loosely described as tax “twins” in practitioner literature.

3.2 Control Theory as a Second Kernel

Following Ivanov, Sethi, Dolgui, and Sokolov's (2018) treatment of supply chain systems as control-theoretic objects, we adopt control theory as a second theoretical kernel alongside digital twin theory. A control-theoretic reading of the Tax Digital Twin treats the enterprise's actual tax posture as a plant subject to continuous disturbance (new transactions, expiring certificates, changing rates, evolving nexus footprint), the twin's virtual entity as a state estimator tracking that plant, and any divergence between estimated and desired (compliant) state as an error signal. This framing does two things for the architecture. First, it supplies a principled basis for the divergence-detection mechanism formalized in Section 4.2, treating divergence not as an ad hoc alert threshold but as the error term in a closed feedback loop. Second, it clarifies the relationship between simulation and control: the twin's forward-simulation capability corresponds to a model-predictive control function, evaluating candidate future actions against the plant model before they are executed in the physical system, rather than reacting only after disturbance has propagated.

3.3 Information Quality as a Boundary Condition

A digital twin's virtual entity is only as trustworthy as the data synchronizing it to physical reality. Wang and Strong's (1996) information quality framework, distinguishing intrinsic, contextual, representational, and accessibility dimensions of data quality, provides the theoretical vocabulary for a boundary condition that recurs throughout this paper: a Tax Digital Twin built on incomplete, stale, or contextually invalid source data will produce confident-looking but unreliable simulations. This is not a minor caveat. Because the twin's principal value proposition is forward simulation rather than retrospective correction, errors in the virtual entity are, by construction, harder to detect than errors in a detection system, since there is no completed transaction against which to validate a projection until the projected future arrives. Section 6 returns to this boundary condition when discussing implementation preconditions.

3.4 Design Science Framing

Consistent with the Design Science Research (DSR) paradigm as articulated by Hevner, March, Park, and Ram (2004) and operationalized through the process model of Peffers, Tuunanen, Rothenberger, and Chatterjee (2007), this paper treats the Tax

Digital Twin as a design artifact — specifically, following March and Smith's artifact taxonomy as applied in the author's prior work, a combination of a model (the formal state representation in Section 4.2) and a method (the reference architecture and its constituent services in Section 4.3). As a conceptual contribution, this paper corresponds to the problem identification, objective definition, and design activities of the DSRM process; the demonstration and evaluation activities — instantiating the architecture against real enterprise data — are identified in Section 6 as the primary direction for subsequent empirical work.

4. The Tax Digital Twin (TDT) Reference Architecture

4.1 Five-Dimension Mapping

Building upon the five-dimension Digital Twin framework proposed by Tao et al. (2018), the Tax Digital Twin (TDT) architecture adapts the core concepts of Physical Entity, Virtual Entity, Twin Data, Services, and Connections to the enterprise tax domain. Together, these five dimensions establish the minimum architectural requirements for continuously representing, synchronizing, simulating, and governing an organization's tax posture.

- **Physical Entity (PE).** The Physical Entity represents the enterprise's operational reality that generates tax obligations. It includes ERP transaction streams, legal entity structures, customer and supplier contracts, inventory movements, warehouse locations, employee deployments, procurement activities, and other business operations that influence nexus, taxability, exemptions, and filing requirements.

- **Virtual Entity (VE).** The Virtual Entity is a continuously maintained computational representation of the enterprise's tax state. Rather than serving as a historical compliance record, it estimates the organization's current tax posture while enabling projection of future tax states under hypothetical business scenarios. The Virtual Entity maintains information such as jurisdiction-specific nexus status, tax registrations, exemption certificate validity, applicable tax rates, filing obligations, and other compliance attributes.

- **Twin Data (TD).** Twin Data provides the synchronization mechanism connecting the Physical and Virtual Entities. It integrates continuously updated information from ERP systems, contract repositories, HR and facilities management systems, logistics platforms, tax engines, and external regulatory data sources. The objective of Twin Data is to ensure that the Virtual Entity remains an accurate and timely representation of the enterprise's operational reality.

- **Services (Ss).** Services expose the analytical capabilities of the Tax Digital Twin to enterprise users and downstream applications. These services include real-time divergence monitoring, what-if scenario simulation, projected tax exposure analysis, audit-readiness assessment, exemption certificate forecasting, and decision support for tax planning.

- **Connections (CN).** Connections comprise the interfaces linking all architectural components. These include APIs connecting ERP platforms, tax engines, external regulatory feeds, document repositories, analytics services, and synchronization services that maintain continuous communication between the Physical Entity, Twin Data, Virtual Entity, and Service layers.

Together, these five dimensions establish the minimum architectural requirements for a Tax Digital Twin. The Physical Entity provides the observable enterprise environment, Twin Data maintains synchronization, the Virtual Entity maintains the enterprise tax state, Services expose predictive capabilities, and Connections enable continuous interaction among these elements.

Table 3. Mapping Tao et al.'s five-dimension Digital Twin framework to enterprise tax compliance.

Digital Twin Dimension	Tax Digital Twin Component	Description
Physical Entity (PE)	Enterprise Operations	ERP transactions, contracts, facilities, inventory, HR
Virtual Entity (VE)	Virtual Tax State	Current and projected enterprise tax posture
Twin Data (TD)	Enterprise & Regulatory Data	ERP, tax engine, HR, contracts, regulatory feeds
Services (Ss)	Predictive Tax Services	Simulation, divergence monitoring, decision support
Connections (CN)	APIs & Synchronization	Interfaces connecting operational and virtual systems

Figure 1. Tax Digital Twin Reference Architecture




Figure 1. Tax Digital Twin (TDT) reference architecture. The architecture integrates enterprise operational systems, external regulatory knowledge, continuous state estimation, and a virtual tax representation to enable predictive tax compliance. Operational data from ERP, HR, contract management, facilities, and tax systems are continuously synchronized to maintain the enterprise's virtual tax state. The Twin Core performs synchronization, divergence analysis, and scenario simulation, while decision support and governance services provide proactive compliance intelligence, auditability, and closed-loop interaction with detection-oriented tax analytics systems.

4.2 Formal State Representation

Let the enterprise's tax state at time t be represented by a multidimensional state vector:

$$S(t)$$

where each element corresponds to a specific jurisdiction and compliance attribute.

Let

J denote the set of jurisdictions in which the enterprise operates, and

K denote the set of tax-state dimensions, including nexus status, applicable tax rate configuration, exemption certificate validity, filing obligations, tax registrations, and other jurisdiction-specific compliance attributes.

Two complementary representations of the enterprise tax state are maintained.

Physical State

$$S^P(t)$$

represents the enterprise's actual tax posture as it exists within operational systems. Because this state cannot be observed directly, it is inferred from ERP transactions, contract repositories, logistics systems, HR records, facilities information, and other operational data sources.

Virtual State

$$S^V(t)$$

represents the Tax Digital Twin's continuously updated estimate of the enterprise's tax posture maintained within the Virtual Entity.

At each synchronization interval, the synchronization engine updates $\mathbf{S}^v(\mathbf{t})$ toward the estimated physical state using continuously arriving enterprise data. The consistency between the two states is evaluated through a divergence metric:

$$D(\mathbf{t}) = \sum_i \in J \sum_k \in K w_k \cdot \text{dist}(\mathbf{S}_{i,k}^p(\mathbf{t}), \mathbf{S}_{i,k}^v(\mathbf{t}))$$

where:

dist(·) denotes an attribute-specific distance function (e.g., binary mismatch for nexus status, normalized numerical difference for tax rates, temporal deviation for certificate expiration dates);

w_k represents the materiality weight assigned to tax-state dimension k ; and

materiality weights are calibrated according to organizational risk tolerance, regulatory significance, and historical compliance impact, consistent with weighted risk prioritization approaches used in prior tax detection research (Kumar, 2026).

When the divergence exceeds a predefined threshold δ , the architecture initiates a reconciliation workflow. Importantly, elevated divergence is informative rather than automatically indicative of non-compliance. Divergence may arise from:

incomplete or stale synchronization data,

delayed operational updates,

recently introduced regulatory changes,

legitimate changes in enterprise operations, or

previously undetected compliance issues.

Accordingly, divergence functions as an early-warning indicator rather than a deterministic error classification.

The defining characteristic of the Tax Digital Twin is its ability to simulate future tax states that have not yet occurred in the physical enterprise. Given a hypothetical business action $\mathbf{a}$, the twin projects a future virtual state according to

$$\mathbf{S}^v(\mathbf{t} + \Delta | \mathbf{a}) = \mathbf{F}(\mathbf{S}^v(\mathbf{t}), \mathbf{a}, \mathbf{R}, \Delta)$$

where:

F denotes the simulation function,

$\mathbf{a}$ represents a hypothetical business action (e.g., opening a warehouse, changing a supply chain, restructuring contracts, entering a new jurisdiction, or migrating to another tax engine),

R denotes the version-controlled regulatory knowledge base containing taxability rules, nexus thresholds, exemption policies, filing requirements, and tax rate schedules, and

Δ represents the simulation horizon.

Unlike static rule-based tax systems, **R** is continuously synchronized with regulatory updates, ensuring that projected tax states remain consistent with the legal framework applicable during the simulated time horizon.

The simulation function produces a projected enterprise tax state together with estimated tax exposure, compliance risk scores, confidence indicators, and recommended mitigation actions. These outputs enable enterprise stakeholders to evaluate alternative operational strategies before implementation, a capability fundamentally unavailable in detection-oriented systems that analyze only completed transactions.

4.3 Layered Architecture

The Tax Digital Twin is organized into six architectural layers that transform operational data into predictive tax intelligence.

Layer 1 — Sensing and Data Ingestion

This layer continuously captures operational signals from ERP transaction streams, procurement systems, contract management platforms, HR systems, facilities databases, tax engines, and external regulatory sources. These inputs collectively provide the observable evidence from which the enterprise's physical tax state can be estimated.

Layer 2 — State Estimation

Rather than constructing the physical state itself, this layer estimates the observable physical tax state $S^p(t)$ by integrating and normalizing heterogeneous enterprise data. Jurisdiction mapping, entity resolution, temporal alignment, and data quality validation are performed to produce a consistent representation of operational reality.

Layer 3 — Twin Core

The Twin Core maintains the Virtual State $S^v(t)$ and continuously synchronizes it with the estimated Physical State. This layer computes the divergence metric $D(t)$, maintains the version-controlled regulatory knowledge base R , performs state synchronization, and provides the computational foundation for predictive simulation.

Layer 4 — Scenario Simulation Services

This layer exposes the simulation engine through standardized APIs, allowing finance, tax, supply chain, and enterprise architecture stakeholders to evaluate hypothetical business decisions. Multiple candidate scenarios may be analyzed and ranked according to projected tax exposure, compliance risk, and financial impact before implementation.

Layer 5 — Decision Support and Visualization

Decision-support services present the current enterprise tax state, divergence alerts, scenario comparisons, projected tax liabilities, and recommended mitigation strategies through interactive dashboards and analytical interfaces. This layer enables proactive tax decision-making rather than retrospective compliance reporting.

Layer 6 — Governance and Audit Trail

The governance layer maintains complete traceability of synchronization events, simulated scenarios, regulatory rule versions, model updates, and decisions influenced by the Digital Twin. Comprehensive versioning and audit logging support organizational governance, regulatory transparency, explainability, and future AI accountability requirements.

Table 4. Functional responsibilities of the six-layer Tax Digital Twin architecture

Layer	Component	Primary Function	Key Output
Layer 1	Sensing & Data Ingestion	Capture enterprise operational data	Raw operational data
Layer 2	State Estimation	Estimate physical tax state $S^p(t)$	Estimated physical state
Layer 3	Twin Core	Synchronization, divergence analysis, rule management	Virtual tax state $S^v(t)$
Layer 4	Scenario Simulation	Evaluate hypothetical business actions	Projected tax state
Layer 5	Decision Support	Visualization and recommendations	Compliance intelligence
Layer 6	Governance & Audit	Version control, explainability, audit logs	Governance artifacts

Figure 2. Six-layer Tax Digital Twin (TDT) architecture.




Figure 2. Six-layer Tax Digital Twin (TDT) architecture. The proposed Tax Digital Twin is organized into six functional layers that transform enterprise operational data into predictive tax intelligence. The architecture progresses from data ingestion and

physical state estimation through the Twin Core, where synchronization, divergence analysis, and regulatory knowledge management are performed, to scenario simulation, decision support, and governance. Together, these layers enable continuous synchronization of enterprise tax posture, predictive evaluation of hypothetical business actions, and closed-loop governance for proactive tax compliance

4.4 Relationship to Detection-Oriented Systems

The Tax Digital Twin is intended to complement rather than replace detection-oriented tax analytics systems such as TaxRadar (Kumar, 2026). The two architectures address different stages of enterprise tax management while reinforcing one another through closed-loop feedback.

The Tax Digital Twin continuously estimates the enterprise's current and future tax states, identifies emerging divergence, and evaluates hypothetical operational changes before implementation. Detection-oriented systems, in contrast, analyze completed transactions to identify anomalies, compliance violations, or tax calculation errors that have already occurred.

The two architectures may therefore be integrated bidirectionally. First, elevated divergence values $D(t)$ generated by the Digital Twin can prioritize jurisdictions, legal entities, or transaction categories for detailed investigation by a detection engine. Second, confirmed findings generated by the detection system may be fed back into the Twin Core to improve state estimation and accelerate synchronization between the Physical and Virtual States.

This complementary relationship reflects the control-theoretic perspective introduced in Section 3. Detection systems primarily operate after operational disturbances have propagated through enterprise processes, whereas the Tax Digital Twin enables organizations to anticipate the tax consequences of future business decisions before those disturbances occur. Together, predictive simulation and post-transaction detection establish a continuous feedback loop that supports proactive, adaptive, and resilient enterprise tax management.

Table 5. Functional comparison between TaxRadar and the proposed Tax Digital Twin.

Capability	TaxRadar	Tax Digital Twin
Detect tax anomalies	✓	Partial
Continuous tax-state estimation	Partial	✓
What-if scenario simulation	X	✓
Predict future tax posture	X	✓
Divergence monitoring	✓	✓
Strategic tax planning	X	✓
Decision support	Limited	Extensive
Integration with ERP	✓	✓
Closed-loop governance	Partial	✓

5. Illustrative Scenarios

The following scenarios illustrate representative enterprise decisions that motivate the proposed Tax Digital Twin (TDT). Unlike detection-oriented systems, which analyze completed transactions after compliance events have occurred, the TDT evaluates hypothetical operational changes before implementation by projecting future enterprise tax states. Each scenario demonstrates how the architecture introduced in Section 4 supports predictive decision-making through state estimation, simulation, and divergence analysis.

5.1 Multi-Jurisdictional Expansion

Initial State.

An enterprise currently operates distribution facilities in five U.S. states and has established nexus only within those jurisdictions.

Hypothetical Action.

Management proposes opening a new distribution center in an additional state to reduce logistics costs and improve delivery times.

Simulation.

The proposed facility is submitted to the simulation function **F** as action **a**. Using the current virtual state **S'(t)**, the version-controlled regulatory knowledge base **R**, projected sales volumes, and existing distribution patterns, the Digital Twin estimates the future tax state **S'(t + Δ | a)**.

Projected Outcomes.

projected nexus creation across affected jurisdictions;
estimated dates for economic nexus threshold attainment;
projected registration and filing obligations;
expected tax liability and compliance costs; and
confidence indicators associated with the projection.

Decision Support.

Decision makers compare alternative facility locations before committing operational resources, enabling tax implications to become an input to strategic planning rather than a consequence discovered after expansion.

5.2 Exemption Certificate Portfolio Decay

Initial State

A multinational B2B enterprise maintains thousands of customer exemption certificates with varying expiration dates across multiple jurisdictions.

Hypothetical Action.

Tax managers evaluate certificate renewal priorities for the next six months.

Simulation.

Certificate validity is maintained as a dynamic state dimension within the Virtual Entity. The simulation projects certificate status over horizon **Δ**, identifies certificates approaching expiration, and estimates the operational and financial consequences if renewals are not completed.

Projected Outcomes.

certificates expected to expire;
projected taxable revenue at risk;
jurisdictions with highest compliance exposure;
expected customer transaction impact; and
prioritized renewal recommendations.

Decision Support.

Rather than responding after transactions fail due to expired documentation, the enterprise proactively schedules renewal activities according to projected business impact.

5.3 Tax Engine Migration

Initial State.

An enterprise plans to migrate from one enterprise tax engine to another while preserving existing tax determination behavior across thousands of transaction scenarios.

Hypothetical Action.

The proposed configuration is introduced as an alternative regulatory rule base **R₂**, while the current production configuration remains represented by **R₁**.

Simulation.

Historical and synthetic transaction sets are executed simultaneously through both rule bases using identical business conditions. The Digital Twin estimates the resulting future tax states under each configuration and computes divergence between the projected outcomes.

Projected Outcomes.

jurisdictions with inconsistent tax calculations;

product taxability differences;
 exemption processing inconsistencies;
 projected financial exposure;
 configuration conflicts requiring correction before deployment.

Decision Support.

The organization validates migration quality before production cutover, significantly reducing the likelihood of widespread compliance errors caused by configuration defects.

5.4 Post-Merger Integration

Initial State.

Following a merger or acquisition, two organizations possess independent legal entities, tax registrations, nexus footprints, exemption certificate repositories, and tax engine configurations.

Hypothetical Action.

A proposed post-merger operating model is submitted to the Digital Twin before ERP and tax system consolidation.

Simulation.

The Virtual Entity combines the tax states of both organizations under a simulated integration scenario. State estimation identifies overlaps, inconsistencies, and newly emerging compliance obligations associated with the merged enterprise.

Projected Outcomes.

consolidated nexus footprint;
 duplicated or conflicting tax registrations;
 exemption certificate overlaps and coverage gaps;
 inconsistent jurisdiction mappings;
 projected filing obligations; and
 estimated compliance costs associated with the integration strategy.

Decision Support.

Tax, finance, and enterprise architecture teams evaluate alternative integration strategies before system consolidation, reducing implementation risk and avoiding costly post-merger remediation.

Table 6. Summary of illustrative Tax Digital Twin scenarios and supported enterprise decisions.

Scenario	Hypothetical Action	Twin Simulation Output	Decision Supported
Multi-Jurisdiction Expansion	Open new distribution center	Nexus projection, filing obligations, compliance cost	Facility location selection
Exemption Certificate Portfolio	Future certificate expiration	Revenue at risk, renewal priorities	Certificate renewal planning
Tax Engine Migration	Alternative rule configuration	Configuration divergence, projected tax exposure	Migration readiness assessment
Post-Merger Integration	Enterprise consolidation	Combined tax state, compliance conflicts	Integration planning

6. Discussion

The Tax Digital Twin (TDT) architecture proposed in this paper contributes to both Digital Twin research and enterprise tax information systems by extending Digital Twin theory into a domain that has received little attention within the existing literature. Whereas most Digital Twin research has focused on manufacturing systems, industrial assets, healthcare, and smart infrastructure, this work conceptualizes enterprise tax compliance as a continuously evolving organizational state that can be synchronized, simulated, and governed through a virtual representation. In doing so, the paper reframes tax management from a retrospective compliance activity into a predictive decision-support capability.

From a theoretical perspective, the proposed architecture integrates three complementary perspectives. Digital Twin theory provides the conceptual foundation for maintaining synchronized physical and virtual representations of enterprise tax posture.

Control theory introduces a closed-loop mechanism in which divergence between physical and virtual states serves as a continuous feedback signal supporting proactive intervention. Information quality theory establishes the boundary conditions under which simulation remains reliable, emphasizing that predictive capability depends fundamentally upon the completeness, timeliness, and contextual validity of operational data. Collectively, these perspectives extend the application of Digital Twin concepts beyond physical assets toward regulatory and governance-intensive enterprise processes.

The architecture also contributes to Design Science Research by introducing a reusable conceptual artifact comprising both a formal state representation and a layered reference architecture. Rather than proposing a software implementation for a particular ERP platform or tax engine, the artifact establishes reusable design principles that may guide future system development across heterogeneous enterprise environments.

From a practical perspective, the architecture provides enterprise tax organizations with a framework for evaluating the tax implications of operational decisions before implementation. Traditional tax analytics primarily identify compliance issues after transactions have occurred. In contrast, the proposed Digital Twin enables organizations to evaluate hypothetical business scenarios—including market expansion, tax engine migration, merger integration, and exemption certificate management—before operational changes are executed. This capability has the potential to support tax planning, enterprise architecture, supply chain design, finance transformation, and strategic investment decisions by incorporating tax consequences directly into organizational planning processes.

Several boundary conditions should nevertheless be acknowledged.

First, the fidelity of the Virtual Entity is fundamentally constrained by the quality, completeness, and timeliness of the data synchronizing it with operational reality, consistent with the information quality framework discussed in Section 3.3 (Wang & Strong, 1996). Incomplete facilities information, inaccurate legal entity structures, delayed transaction synchronization, or inconsistent jurisdiction mappings may all reduce the accuracy of projected tax states. Because future simulations cannot be validated immediately, maintaining robust enterprise data governance becomes a prerequisite for reliable Digital Twin operation rather than a secondary implementation consideration.

Second, maintaining a continuously synchronized and simulable tax state across a multinational enterprise introduces significant computational and governance requirements. The regulatory knowledge base **R** must continuously reflect evolving tax legislation, nexus thresholds, filing obligations, exemption rules, and tax rate schedules across multiple jurisdictions. To minimize maintenance overhead and avoid regulatory inconsistencies, future implementations should integrate the Digital Twin with the enterprise's production tax engine through a shared, version-controlled regulatory knowledge repository rather than maintaining duplicate rule bases.

Third, the proposed architecture should be viewed as complementary to detection-oriented tax analytics rather than as a replacement. Detection systems analyze completed transactions to identify compliance anomalies, whereas the Tax Digital Twin projects future organizational tax states before operational decisions are executed. Together these capabilities establish a continuous feedback loop in which predictive simulation informs business decisions, while post-transaction detection continuously improves state estimation and model accuracy.

Finally, this paper represents a conceptual Design Science artifact and therefore does not include empirical evaluation. Although the architecture is grounded in established Digital Twin theory, control theory, and information quality research, its practical effectiveness remains to be demonstrated through implementation within realistic enterprise environments. Accordingly, the findings should be interpreted as design propositions rather than empirically validated performance claims.

Future research should focus on instantiating the proposed architecture within enterprise ERP and tax technology ecosystems to evaluate synchronization accuracy, simulation fidelity, computational scalability, and organizational value. Additional opportunities include integrating machine learning models to improve state estimation, employing AI agents to automate scenario generation and recommendation, investigating reinforcement learning approaches for adaptive tax planning, and extending the architecture toward multi-enterprise Digital Twins capable of modeling supply-chain-wide tax interactions. Such investigations would move the Tax Digital Twin from a conceptual reference architecture toward an operational platform for predictive enterprise tax management.

7. Conclusion

Enterprise taxation has traditionally been managed through retrospective compliance processes that identify calculation errors, filing deficiencies, and regulatory violations only after business transactions have occurred. While recent advances in AI-based

tax analytics have significantly improved anomaly detection and compliance monitoring, these approaches remain fundamentally reactive because they operate on historical operational data. This paper argues that predictive enterprise tax management requires a different architectural paradigm—one capable of representing and simulating an organization's evolving tax posture before operational decisions are executed.

To address this challenge, the paper introduced the **Tax Digital Twin (TDT)**, a conceptual reference architecture that extends Digital Twin theory from its traditional application in manufacturing and industrial systems to the domain of enterprise tax compliance. Grounded in the Physical–Virtual Twin model of Grieves and Vickers (2017), enriched by a control-theoretic perspective on continuous feedback, and positioned within the Design Science Research paradigm, the proposed architecture formalizes enterprise tax posture as a continuously synchronized virtual state capable of forward simulation. In doing so, the paper contributes a reusable conceptual artifact comprising a formal state representation, a divergence-based synchronization mechanism, a six-layer reference architecture, and a simulation framework for evaluating hypothetical business actions before implementation.

The proposed architecture also clarifies the complementary relationship between predictive simulation and detection-oriented tax analytics. Rather than replacing AI-based tax detection systems, the Tax Digital Twin extends their capabilities by enabling organizations to anticipate the tax consequences of future operational decisions while continuously incorporating feedback from observed compliance outcomes. This closed-loop interaction establishes a more proactive model of enterprise tax governance in which simulation and detection reinforce one another throughout the organizational decision-making process.

The illustrative scenarios demonstrate how the proposed architecture supports decisions that cannot be addressed by detection-oriented systems alone, including multi-jurisdictional expansion, exemption certificate portfolio management, tax engine migration, and post-merger integration. More broadly, the architecture illustrates how Digital Twin principles can be applied to regulatory and governance-intensive enterprise domains whose primary challenge is not physical asset monitoring but continuous organizational state management under evolving legal and business conditions.

As a conceptual Design Science artifact, the Tax Digital Twin is intended to provide a foundation for future research rather than a completed operational system. The next stage of this research program is the implementation and empirical evaluation of the architecture within enterprise ERP and tax technology environments to assess synchronization fidelity, simulation accuracy, computational scalability, organizational impact, and decision quality. Future investigations may also explore the integration of AI agents, machine learning, and adaptive optimization techniques to enhance autonomous tax-state estimation and scenario generation.

Ultimately, this work suggests that the future of enterprise tax technology extends beyond detecting compliance failures toward continuously anticipating them. By introducing the concept of a Tax Digital Twin, this paper seeks to establish a new research direction at the intersection of Digital Twin technology, enterprise information systems, artificial intelligence, and regulatory compliance, providing both researchers and practitioners with a foundation for predictive, adaptive, and continuously governed enterprise tax management.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

- [1] **Barricelli, B. R., Casiraghi, E., & Fogli, D. (2019).** A survey on digital twin: Definitions, characteristics, applications, and design implications. *IEEE Access*, 7, 167653–167671. <https://doi.org/10.1109/ACCESS.2019.2953499>
- [2] **Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020).** Digital twin: Enabling technologies, challenges and open research. *IEEE Access*, 8, 108952–108971. <https://doi.org/10.1109/ACCESS.2020.2998358>
- [3] **Glaessgen, E. H., & Stargel, D. S. (2012).** *The digital twin paradigm for future NASA and U.S. Air Force vehicles*. In 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference. American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2012-1818>
- [4] **Grieves, M. (2014).** *Digital twin: Manufacturing excellence through virtual factory replication* (White paper, pp. 1–7). Michael W. Grieves, LLC.

- [5] **Grieves, M., & Vickers, J. (2017).** Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), *Transdisciplinary Perspectives on Complex Systems* (pp. 85–113). Springer. https://doi.org/10.1007/978-3-319-38756-7_4
- [6] **Gu, X., Tian, X., Yang, J., & Liu, M. (2025).** Building and performance validation of a digital twin regulatory framework for financial compliance and market transparency. Research Square. <https://doi.org/10.21203/rs.3.rs-8220545/v1>
- [7] **Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004).** Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. <https://doi.org/10.2307/25148625>
- [8] **Ivanov, D., & Dolgui, A. (2021).** A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0. *Production Planning & Control*, 32(9), 775–788. <https://doi.org/10.1080/09537287.2020.1768450>
- [9] **Ivanov, D., Dolgui, A., Das, A., & Sokolov, B. (2019).** Digital supply chain twins: Managing the ripple effect, resilience, and disruption risks by data-driven optimization, simulation, and visibility. In A. Ivanov, A. Dolgui, & B. Sokolov (Eds.), *Handbook of Ripple Effects in the Supply Chain* (pp. 309–332). Springer
- [10] **Ivanov, D., Sethi, S. P., Dolgui, A., & Sokolov, B. (2018).** A survey on control theory applications to operational systems, supply chain management, and Industry 4.0. *Annual Reviews in Control*, 46, 134–147.
- [11] **Jones, D., Snider, C., Nassehi, A., Yon, J., & Hicks, B. (2020).** Characterising the Digital Twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology*, 29(Part A), 36–52. <https://doi.org/10.1016/j.cirpj.2020.02.002>
- [12] **Kritzinger, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018).** Digital twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine*, 51(11), 1016–1022. <https://doi.org/10.1016/j.ifacol.2018.08.474>.
- [13] **OECD. (2020).** *Tax Administration 3.0: The Digital Transformation of Tax Administration*. OECD Publishing, Paris. <https://doi.org/10.1787/ca274cc5-en>
- [14] **Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007).** A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77. <https://doi.org/10.2753/MIS0742-1222240302>
- [15] **Qi, Q., Tao, F., Hu, T., Anwer, N., Liu, A., Wei, Y., Wang, L., & Nee, A. Y. C. (2021).** Enabling technologies and tools for digital twin. *Journal of Manufacturing Systems*, 58(Part B), 3–21. <https://doi.org/10.1016/j.jmsy.2019.10.001>
- [16] **South Dakota v. Wayfair, Inc., 585 U.S. 162 (2018)**
- [17] **Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019).** Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415. <https://doi.org/10.1109/TII.2018.2873186>
- [18] **Wang, R. Y., & Strong, D. M. (1996).** Beyond accuracy: What data quality means to data consumers. *Journal of Management Information Systems*, 12(4), 5–33. <https://doi.org/10.1080/07421222.1996.11518099>
- [19] **Williamson, O. E. (1985).** *The economic institutions of capitalism: Firms, markets, relational contracting*. Free Press
- [20] **Xu, C., Bai, T., Zhou, C., & Huang, S. (2025).** Digital tax administration and enterprise innovation: Evidence from the Golden Tax Project III in China. *Research in International Business and Finance*, 77(Part A). <https://doi.org/10.1016/j.ribaf.2025.102829>
- [21] **Zhang, N., Bahsoon, R., Tziritas, N., & Theodoropoulos, G. (2024).** Knowledge equivalence in digital twins of intelligent systems. *ACM Transactions on Modeling and Computer Simulation*, 34(1), 3:1–3:37