Computer Fraud and Security
ISSN (online): 1873-7056

Enterprise Tax Systems Modernization Using Intelligent Integration Frameworks

Ajay Mirani

Independent Researcher, USA

Abstract: Corporate tax departments continue to operate on spreadsheets and manual review cycles even as adjacent finance domains have modernized around real-time data and machine learning. This paper introduces the Intelligent Tax Integration Framework (ITIF), a six-layer domain-specific architecture for artificial intelligence (AI)-driven enterprise tax systems modernization. ITIF addresses six interlocking challenges: fragmented multi-jurisdiction data, tax classification at the stock-keeping unit (SKU) scale, real-time transfer pricing (TP) monitoring, indirect tax determination for ambiguous categories, current and deferred tax provision under Accounting Standards Codification (ASC) 740 and International Accounting Standard (IAS) 12, and filing readiness across an expanding patchwork of digital reporting regimes. Each layer pairs a specific AI technique—fine-tuned Bidirectional Encoder Representations from Transformers (BERT), an Isolation Forest and long short-term memory (LSTM) ensemble with Shapley additive explanations (SHAP) attribution, Monte Carlo simulation for uncertain tax positions, and Retrieval-Augmented Generation (RAG) for audit queries—with regulatory knowledge graphs anchored to Organization for Economic Co-operation and Development (OECD), Financial Accounting Standards Board (FASB), and International Accounting Standards Board (IASB) source material. A simulated case study illustrates deployment across eighteen legal entities and twelve jurisdictions, grounded in practitioner experience from a Systems, Applications, and Products (SAP) to Microsoft Dynamics AX migration at a high-volume global technology manufacturer. A nine-dimension comparative evaluation positions ITIF against legacy tax systems and rule-based tax engines and identifies continuous tax intelligence coverage during enterprise resource planning (ERP) cutover as the framework's most distinctive capability.

Keywords: *enterprise tax systems, transfer pricing, indirect tax automation, tax provisions, ASC 740, BEPS compliance, ERP migration, explainable AI*

I. INTRODUCTION

Enterprise tax systems occupy a peculiar position in the corporate technology stack. They are simultaneously among the most consequential applications an enterprise runs and among the most architecturally antiquated. Tax restatements rival financial reporting restatements in magnitude, and egregious failures carry criminal liability; yet a practitioner from 2005 would recognize the operating model at most large enterprises today. Provision is still driven out of spreadsheets. Indirect tax determination is still reviewed manually at the margins. Transfer pricing is still outsourced to consultants on quarterly cycles. The institutional memory of a handful of specialists still determines the effectiveness of audit defense.

The underinvestment is not simply a matter of budget or organizational attention. Tax intelligence is structurally cross-cutting: transfer pricing reaches into every intercompany transaction, indirect tax determination fires on every sale, and provision depends on the complete financial position of every legal entity in the group. A tax system designed as a standalone application—one that ingests nightly batch exports from the ERP and returns determinations to be posted by hand—cannot participate in the real-time, continuous governance that the rest of the finance stack increasingly demands.

The Intelligent Tax Integration Framework (ITIF) is a domain-specific AI architecture that responds to that mismatch directly. Rather than layering new capabilities on an old integration pattern, ITIF is designed from inception as a tax-intelligence tier over a modern event-driven financial platform. This paper is the third in a series on AI-driven enterprise financial governance. Paper 1 introduced the Adaptive Compliance Intelligence and Governance (ACIG) framework [1]—a six-layer architecture for continuous, AI-driven compliance intelligence. Paper 2 extended this framework to full AI-native platform design [2], situating ACIG as the compliance plane of a broader intelligent financial platform. This paper applies domain-specific tax intelligence architecture to the enterprise tax function, extending the series into a third discipline. Its design choices draw on two bodies of prior work: the AI-native enterprise financial governance framework established by Papers 1 and 2 of this series [1][2] and the literature documenting the practical challenges of deploying machine learning in production enterprise systems [3].

The framework is grounded in practice. Many of its defining features—parallel-run support during ERP migration, SHAP calibration to TP documentation standards, Monte Carlo treatment of uncertain tax positions—trace directly to challenges encountered during multi-year SAP-to-Microsoft Dynamics AX transitions at a high-volume global technology manufacturer. Section II enumerates the contributions. Sections III and IV situate the work and list the six tax intelligence challenges the framework addresses. Section V specifies the architecture. Section VI defers integration with prior work until after the architecture has been fully specified so that comparisons can be drawn against a defined framework rather than a preview; this is a deliberate structural choice. Section VII proposes a maturity model. Sections VIII and IX present a case study and comparative evaluation. Sections X and XI close with discussion and conclusion.

II. PRIMARY CONTRIBUTIONS

This paper makes five contributions to enterprise AI, tax technology, and financial governance.

First, the ITIF itself: a six-layer architecture that treats tax intelligence as an integrated domain rather than as six disconnected point problems. The framework spans data unification through classification, transfer pricing, indirect tax, provision, and filing. To our knowledge, no prior published architecture addresses enterprise tax as a unified AI problem at this level of structural specificity, and none pairs AI technique selection with the regulatory knowledge graphs required to make output admissible in examination.

Second, a transfer pricing intelligence module is described that specifies an Isolation Forest and LSTM ensemble anchored to the OECD Transfer Pricing Guidelines knowledge graph [7], with SHAP-based feature attribution calibrated to the documentation requirements of major tax authorities. Transfer pricing is treated here as a first-class machine learning problem, not as a consulting engagement. The SHAP calibration, in particular, is a design choice without precedent in existing TP monitoring tools and is what makes the module's output usable directly as contemporaneous documentation.

Third, a provision-intelligence architecture that uses Monte Carlo simulation for uncertain tax position (UTP) assessment under ASC 740 [9] and IAS 12 [10], along with an ML-based reconciliation model for provision-to-return variance analysis. This targets the highest-risk and least-automated segment of the tax close cycle, where the gap between estimation judgment and realized outcome is typically the largest and where tax directors currently operate with the least decision support.

Fourth, a five-level Tax Systems Modernization Maturity Model (TSMMM) that gives enterprises a staged adoption path calibrated to organizational realities rather than technological aspiration. The model is designed to surface the likely binding constraint at any given stage—data quality, model governance, or organizational readiness—rather than simply prescribe what capability to build next.

Fifth, a simulated multi-jurisdiction ERP migration case study that illustrates how ITIF addresses the tax continuity gap between legacy cutover and new-system stabilization—a period when standalone tax applications lose coverage and during which compliance risk compounds quickly without an obvious detection surface. The case is grounded in practitioner experience at a high-volume global technology manufacturer.

III. BACKGROUND AND RELATED WORK

A. The Enterprise Tax Technology Landscape

Specialized vendors—Thomson Reuters ONESOURCE, Vertex, Avalara, and Oracle Tax Reporting Cloud—dominate enterprise tax technology. Each is comprehensive within its domain, and each is architected as a standalone application fed by batch exports. That model is the source of the integration cost documented across enterprise ML deployments [2]: every new analytical capability requires bespoke pipeline work that grows faster than the platform. A Deloitte survey found that 67 percent of Fortune 500 tax departments still run provision calculations primarily through spreadsheets, and 54 percent cite data access latency as the primary barrier to effective tax risk management [21].

B. AI Applications in Tax

Research on AI in tax has addressed isolated sub-problems. Ash, Guillot, and Han applied text classifiers to large corpora of tax legislation to automatically tag and categorize tax law provisions across jurisdictions, demonstrating that supervised NLP scales to statutory tax text [4]. Weber et al. applied graph convolutional networks to large-scale financial transaction graphs for anomaly detection in anti-money-laundering contexts [5]; that technique is a direct analog to the transfer pricing layer developed here, where intercompany chains form a similar graph substrate. Chalkidis et al. showed that domain-adapted transformer models consistently perform better than generic language models on legal classification tasks [6]. This same pattern motivates the hybrid rule-plus-ML override design of Layer 4, where a domain-adapted model handles the ambiguous long tail that rule engines classify poorly. What has been missing from the literature is an integrated framework that binds these techniques into a coherent tax-domain architecture. ITIF fills that gap.

C. Transfer Pricing and BEPS

Transfer pricing is now the highest-profile tax risk for most multinationals. The OECD Base Erosion and Profit Shifting (BEPS) project [15] and its Pillar Two minimum tax frameworks have intensified scrutiny, while the arm's length principle requires pricing consistency across hundreds of thousands of intercompany transactions per day at a large group. Monitoring that standard in real time is beyond any manual review process. Recent work on AI-driven BEPS compliance monitoring [17] and real-time intercompany analytics [18] demonstrates the feasibility of the ensemble approach adopted in Layer 3. ITIF's Layer 3 addresses the gap directly through the ensemble architecture specified in Section V.

D. Prior Work on Compliance Architecture and ML Systems

ITIF draws on two research traditions. The first is the AI-native enterprise financial governance framework developed in Papers 1 and 2 of this series: the ACIG compliance intelligence architecture [1] and the AI-native platform patterns that host it [2], together with the broader literature on knowledge graph construction that underpins regulatory compliance mapping [16]. The second is the study of architectural patterns required to deploy machine learning at enterprise scale [3], including guidance on managing model risk from banking supervisors [14]. On transfer pricing specifically, this paper extends the monitoring approaches surveyed in the TP automation literature, including work on AI-driven BEPS compliance monitoring [17] and the growing practitioner literature on real-time intercompany analytics [18]. ITIF is a tax-domain application of principles established in both traditions, not a clean-sheet construction.

E. ERP Migration and Tax Continuity

The problem of maintaining tax compliance across an ERP migration has received limited academic attention relative to its operational weight. Coşkun et al. systematically mapped 72 peer-reviewed studies of ERP failure and catalogued recurring risk factors across integration debt, change management, and post-cutover operational stability [8]; tax and regulatory continuity sit squarely inside the integration-risk cluster they identify. That mapping directly motivates the parallel-run migration support capability detailed in the case study and is consistent with practitioner experience: the migration window is typically the single highest-risk period in a multi-year tax technology calendar.

IV. FOUNDATIONAL TAX INTELLIGENCE CHALLENGES

Six systemic challenges define the layers of ITIF.

Challenge 1: Tax Data Fragmentation

Tax functions consume data from dozens of source systems—ERP, billing, intercompany accounting, fixed assets, payroll, and treasury—that use different data models, currencies, entity hierarchies, and period definitions. Assembling a reconciled view of the enterprise tax position is a manual exercise that consumes a substantial share of the team's capacity during close. Multi-ERP environments, common after acquisition-driven growth or multi-region rollout, compound the problem with inconsistent charts of accounts and overlapping entity codes that demand manual normalization before any tax calculation can begin. The cost surfaces most visibly at quarter-end, when the first five to seven days of the cycle are spent preparing data rather than analyzing it.

Challenge 2: Tax Classification at Scale

Every product, service, and transaction type carries a classification that drives indirect tax, income tax, and customs treatment. At a large technology manufacturer with thousands of SKUs sold into dozens of jurisdictions, keeping classifications current is a continuous expert-labor task and one that scales poorly as the product catalog grows. Misclassifications produce both overpayment risk (tax paid that was not due) and underpayment risk, with the latter carrying penalty and interest exposure that can compound across periods before it is detected.

Challenge 3: Transfer Pricing Monitoring

Arm's length pricing must be monitored across every intercompany transaction—procurement, royalties, services, financing, and inventory transfers. In large multinationals, that volume runs into the hundreds of thousands per day. Current practice relies on quarterly or annual TP documentation reviews conducted by external consultants, producing detection windows of months during which violations accumulate before being identified. By the time an issue surfaces, the remediation cost—often including retrospective adjustments, penalty exposure, and advance pricing agreement rework—typically exceeds what continuous monitoring would have cost.

Challenge 4: Indirect Tax Complexity

Value-added tax (VAT), goods and services tax (GST), and sales tax regimes differ sharply across jurisdictions in rates, exemptions, recovery rules, filing frequency, and digital reporting mandates. Indirect tax determinations fire millions of times per day across purchasing, sales, and intercompany flows. Rule engines handle the bulk of cases accurately, but the long tail of edge cases—cross-border transactions, exempt entities, ambiguous product categories, and composite supplies—generates a volume of incorrect determinations that require manual review and, increasingly, defensible documentation.

Challenge 5: Tax Provision Complexity

Computing current and deferred tax provisions under ASC 740 [9] or IAS 12 [10] requires integration of data from every legal entity, application of complex temporary difference calculations, and judgment on uncertain tax positions. The provision cycle routinely consumes two to three weeks of close, with material estimation risk that persists until provision-to-return reconciliation occurs months later. Under time pressure, tax directors typically default to binary more-likely-than-not judgments on UTPs, a posture that offers limited information to governance committees and little auditable rigor when examinations arrive.

Challenge 6: Regulatory Filing and Audit Defense

Documentation demands have sharply increased due to country-by-country reporting, real-time digital mandates (Standard Audit File for Tax, e-invoicing, Pan-European Public Procurement On-Line), and enhanced audit capability at tax authorities. Authorities in major jurisdictions now operate sophisticated data analytics capable of identifying patterns that may indicate aggressive planning or compliance failure. Enterprises that cannot assemble complete, consistent documentation across jurisdictions and periods face disproportionate audit risk and an escalating compliance burden.

V. THE INTELLIGENT TAX INTEGRATION FRAMEWORK

ITIF addresses the six challenges enumerated in Section IV through a six-layer architecture. Table I summarizes the layers. Fig. 1 shows the reference architecture. Fig. 2 traces a representative transaction through all six layers.

TABLE I. ITIF LAYER SUMMARY

ITIF Layer	Primary Purpose	AI / ML Technique	Key Output
L1 Data Unification	Multi-entity, multi-currency aggregation	NLP schema inference, ETL automation	Unified Tax Fabric
L2 Tax Classification	Product / service tax categorization	Fine-tuned BERT, hierarchical classifier	Tax category + confidence

ITIF Layer	Primary Purpose	AI / ML Technique	Key Output
L3 Transfer Pricing	TP arm's length analysis, BEPS compliance	Isolation Forest, LSTM, OECD knowledge graph	TP risk score + SHAP
L4 Indirect Tax	VAT, GST, sales tax determination	Rule engine + ML override, jurisdiction graph	Tax determination
L5 Tax Provision	Current and deferred tax, ETR, ASC 740	Monte Carlo, scenario simulation, reconciliation, ML	Provision report
L6 Filing Intelligence	Filing readiness, cross-jurisdiction consistency	RAG audit query, XGBoost penalty modelling	Filing package

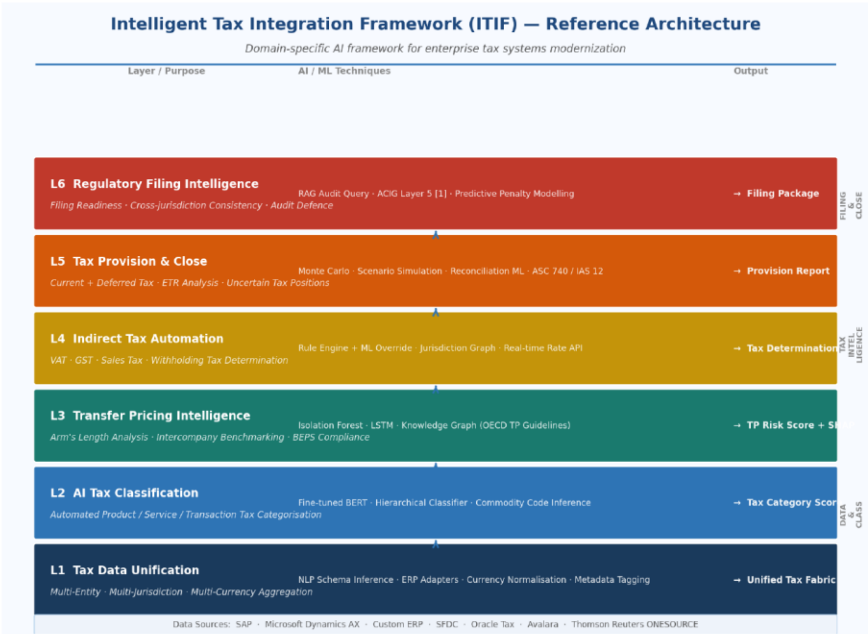

Fig. 1. ITIF Reference Architecture — six-layer tax intelligence stack showing AI techniques, data sources, and outputs for each layer.

Layer 1: Tax Data Unification

The unification layer aggregates financial events from ERP systems (SAP, Dynamics AX, custom platforms), billing, intercompany accounting, fixed asset registers, and treasury, and normalizes them into a unified tax fabric. What distinguishes it from conventional extract-transform-load (ETL) is NLP-driven schema inference across heterogeneous data models, together with entity hierarchy normalization that maps legal entity structures across systems never designed to share a chart of accounts. A currency normalization layer converts all transaction values to functional currency using historical rates, with audit trail entries recording every conversion. That step is a prerequisite for both Layer 3 and Layer 5.

Layer 2: AI Tax Classification Engine

A fine-tuned BERT classifier determines the tax category of every product, service, and transaction type. The classifier is hierarchical: a top-level head assigns broad categories (taxable, exempt, zero-rated, reduced-rate), while jurisdiction-specific sub-classifiers apply detailed rate and exemption rules. A specialized head trained on the tariff schedule corpus infers Harmonized System (HS) and Combined Nomenclature (CN) codes for customs and indirect taxes.

Classification confidence below a configurable threshold triggers human-in-the-loop review, preventing low-confidence determinations from reaching high-value transactions without specialist oversight.

Layer 3: Transfer Pricing Intelligence Module

Layer 3 is the most technically involved layer in the framework. It runs an ML ensemble for real-time intercompany transaction monitoring, anchored to a knowledge graph representation of the OECD Transfer Pricing Guidelines [7] and the enterprise's approved TP policies. Three complementary models contribute: an Isolation Forest over transaction features (amount, entity pair, type, geographic pair, timing) producing score s_IF; an LSTM trained on intercompany time-series producing s_LSTM; and a Benchmarking Comparator that retrieves the arm's length range from the TP knowledge graph and computes a deviation score s_ALP. The composite TP risk score is $S\_TP = 0.30 \cdot s\_IF + 0.35 \cdot s\_LSTM + 0.35 \cdot s\_ALP$, weighted toward the arm's length deviation component to reflect its regulatory significance. Every flag carries a SHAP attribution report [11] calibrated to TP documentation standards.

Layer 4: Indirect Tax Automation

Layer 4 is hybrid by design. A rule-based determination engine handles the bulk of straightforward cases; an ML override intercepts determinations where the rule engine's confidence is low and substitutes a learned prediction, applying the domain-adapted transformer pattern reported by Chalkidis et al. [6] to the ambiguous long tail that rule engines classify poorly. Operationally, "low confidence" is defined as a rule engine confidence score below a configurable threshold θ (default 0.75), estimated as P(correct determination) on a held-out validation set. Determinations with confidence in the range $[\theta, 0.90)$ are routed to the ML override; those below θ are escalated for human review. This three-tier routing — rule engine, ML override, human — mirrors the pattern established in Layer 2 and provides a consistent governance posture across the framework. A jurisdiction knowledge graph maintains a versioned representation of indirect tax rates, exemptions, and recovery rules across all relevant jurisdictions, with update proposals generated automatically from tax authority announcements using the knowledge graph construction techniques surveyed in the broader literature [16]. Real-time rate application programming interface (API) integration handles volatile regimes—notably United States sales tax, where sub-state rates change frequently and graph maintenance alone cannot keep up.

Layer 5: Tax Provision and Close Intelligence

The provision layer implements a calculation engine for current and deferred tax under ASC 740 [9] and IAS 12 [10], consuming the normalized data from Layer 1 and classification output from Layer 2 to compute entity-level provisions that roll up to the consolidated group. The distinctive AI innovation is the Monte Carlo treatment of uncertain tax positions. Each significant UTP is modeled as a probability distribution over outcomes, and simulation generates the expected value and variance of the reserve. That replaces the current practice of binary more-likely-than-not judgment under time pressure with a quantitatively grounded distribution. An ML reconciliation model trained on historical provision-to-return data automates variance analysis.

Layer 6: Regulatory Filing Intelligence

Layer 6 maintains an immutable, cross-jurisdiction-consistent record of tax determinations, provision calculations, and filing positions. A RAG-based query interface [13] lets the tax team respond to regulatory inquiries in natural language against that record. Filing readiness scoring flags gaps and inconsistencies before deadlines. Cross-jurisdiction consistency checks verify that positions taken in one jurisdiction align with related positions elsewhere—critical for country-by-country reporting. An XGBoost penalty model assigns examination-risk scores to filing positions, prioritizing supporting-documentation effort.

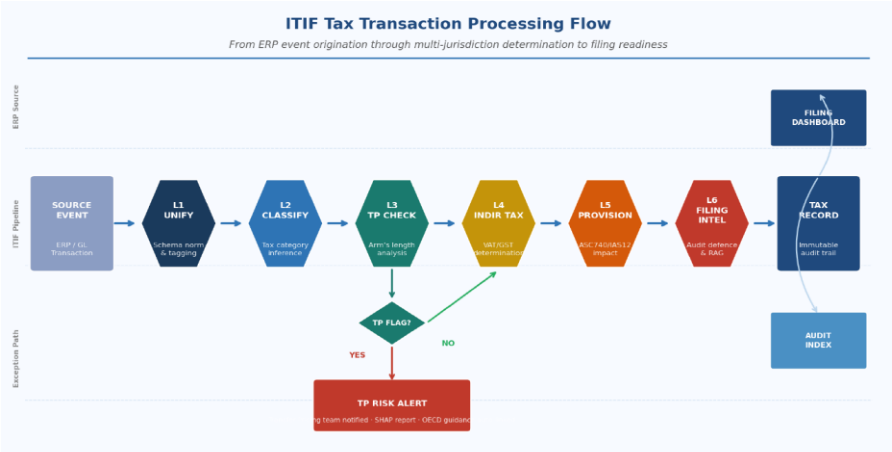

Fig. 2. ITIF Transaction Processing Flow — a representative transaction traversing all six layers from ERP origination through TP decision routing to audit trail indexing.

VI. INTEGRATION WITH PRIOR WORK

A. Relationship to Compliance Intelligence Research

ITIF operationalizes several themes from the machine-learning-for-financial-compliance literature. The anomaly detection ensemble in Layer 3 is a tax-domain adaptation of the ML-driven risk-monitoring architecture established in the ACIG framework [1]. It includes transfer pricing features, regulatory grounding against OECD guidelines [7], and SHAP attribution [11] to meet documentation standards that risk-only systems do not address. Layer 5's provision intelligence generalizes the predictive-governance strand of the same literature into a provision-specific setting, with Monte Carlo UTP simulation supplying probabilistic input to reserve decisions rather than the binary judgment that currently dominates practice. Layer 6 uses general-purpose knowledge graph construction techniques [16] to provide RAG audit queries over an immutable tax-record corpus, changing compliance documentation from a retrospective assembly exercise into a near-real-time retrieval problem.

B. Relationship to ML Systems Architecture

ITIF is also a domain-specific application of enterprise machine learning deployment patterns [2]. Layers 3 and 4 depend on event-driven data access—transfer pricing and indirect tax determination both require transaction visibility within seconds of origination to be operationally useful at scale. The governance regimes that model risk guidance [14] require the Layer 2 classifier and the Layer 3 ensemble to follow production ML standards, which include documented training procedures, ongoing performance monitoring, and reproducible inference. The framework's schema adaptability during ERP migration is an application of the deployment-time schema for resilience documented in ML engineering research [2]. The explainability-first treatment of all ML outputs—every determination and flag carries a SHAP attribution report [11, 12]—is a tax-domain instantiation of that same deployment tradition and is calibrated to meet both internal governance and external regulatory documentation standards.

VII. TAX SYSTEMS MODERNIZATION MATURITY MODEL

The Tax Systems Modernization Maturity Model (TSMMM) is a five-level staged adoption pathway for tax technology transformation. It is a companion tool to ITIF, intended to help enterprises identify their starting point and sequence investments. Each level is characterized not only by the technologies in production but also by the binding constraint that must be relieved before the next level becomes reachable—typically data quality at Levels 1–2, model governance at Levels 3–4, and organizational readiness at Level 5.

Level 1: Manual / Spreadsheet-Dependent

Provision in spreadsheets, manual indirect tax review, external consultants for TP documentation, and audit defense depend on manual documentation assembly. This is the current state for the majority of large enterprise tax departments

[3]. The binding constraint is data quality: tax-relevant data lives in ERP extracts, side spreadsheets, and email attachments that are reconciled by hand at close. Movement to Level 2 requires investment in canonical tax data definitions and source-of-truth discipline rather than net-new software.

Level 2: Structured with Point Solutions

Specialist tax applications—ONESOURCE, Vertex, Avalara, and a TP documentation tool—deployed as standalone receivers of batch data. The batch cycle limits data freshness, and new analytics require custom pipelines. We continue to monitor TP periodically. Organizations at this level often mistake vendor acquisition for transformation; the architecture is still batch-feed-and-point-solution, and ML remains absent from the determination path. The binding constraint shifts from data quality toward integration debt across the point-solution portfolio.

Level 3: AI-Augmented (Partial ITIF)

Data unification and classification (Layers 1–2) in production; ML-assisted indirect tax determination (Layer 4) in early deployment. Real-time transaction data available through early event-driven platform adoption. TP monitoring is still primarily batch but with ML-assisted anomaly flagging. Model governance emerges as the new binding constraint—organizations must stand up documentation. monitoring, and review processes that satisfy both internal audit and external examination before they can extend ML into higher-risk determinations.

Level 4: Intelligent (Full ITIF without Autonomous Filing)

All six layers are in production. Real-time TP monitoring, AI-driven indirect tax, ML-assisted provision, RAG audit query, and cross-jurisdiction consistency checking are operational. Human oversight retained for uncertain positions, material TP decisions, and filing sign-off. The tax function at this level is organized around reviewing exceptions and governance signals rather than preparing returns; headcount composition shifts toward analytic and review skills, and the relationship with external advisors narrows to advisory work on novel positions.

Level 5: Autonomous Tax Intelligence

Routine indirect tax determinations, straightforward TP validations, and settled-position provision calculations processed without human intervention. Human attention concentrated on novel or high-risk positions and strategic planning. Audit responses for routine inquiries generated automatically via Layer 6. Few if any enterprises sit at this level today; it is aspirational. The binding constraint at this stage is organizational and regulatory rather than technical—tax authorities and internal risk committees must develop comfort with automated determinations before a meaningful share of filings can move to autonomous processing.

VIII. ILLUSTRATIVE CASE STUDY: MULTI-JURISDICTION ERP MIGRATION TAX CONTINUITY

The following case study illustrates ITIF in practice. It is a conceptual scenario grounded in practitioner experience managing tax close during ERP transitions at a global technology manufacturer; it is not an empirical result. The organization, transaction volumes, and timings have been generalized to describe a pattern that recurs across multinational enterprises undergoing multi-jurisdiction ERP cutovers. The intent is not to establish measurable performance claims but to show how the six layers interact under realistic operational stress—heterogeneous source systems, regulatory deadlines that do not pause for IT programs, and transfer pricing positions that must remain defensible across the transition.

TABLE II. ILLUSTRATIVE CASE STUDY: ITIF IN ERP MIGRATION TAX CONTINUITY

Organization	GlobalTech Corp (simulated) — multinational technology manufacturer migrating from SAP ECC to Microsoft Dynamics AX across 18 legal entities in 12 jurisdictions over a 24-month program. Annual intercompany volume is approximately 3.2 million transactions across procurement, royalties, and services.
Tax Risk at Migration	A 6-month window of heightened risk: SAP tax determination rules not yet fully replicated in Dynamics AX, entity hierarchies differ between the two

	systems, and TP monitoring previously performed by consultants on SAP exports has no agreed-upon equivalent in the new platform.
ITIF Deployment	ITIF runs in parallel mode alongside both SAP and Dynamics AX from month one. The schema adaptation layer normalizes both systems' output into the unified tax fabric, enabling continuous TP monitoring and indirect tax validation across both systems at the same time. ITIF becomes the single source of truth for the tax position during the transition.
Transfer Pricing Continuity	Layer 3 monitors all 3.2 million annual intercompany transactions in real time across both ERPs. In month three, the TP module flags 847 APAC service-fee transactions below the documented arm's length range. (The alert threshold is set at the 97th percentile of the S_TP distribution computed over a 12-month baseline window, consistent with the ensemble specification in Section V—approximately 3 percent of all intercompany transactions are escalated for human review.) Composite score S_TP = 0.87 triggers an alert with SHAP attribution identifying s_ALP = 0.91 as the primary driver. The transactions trace to a pricing table that did not migrate correctly during Dynamics AX cutover; the issue is remediated before the quarter-end TP documentation deadline.
Provision Close Acceleration	Layer 5 eliminates the manual data assembly and reconciliation that previously consumed the first week of the provision cycle. With normalized, reconciled entity-level data available on demand, provision calculation begins immediately at period close. Monte Carlo simulation of three material UTPs completes in roughly four hours, yielding a quantitative reserve estimate in place of the prior binary judgment under time pressure.
Audit Defense Readiness	When a tax authority issues an information request on intercompany royalty arrangements, Layer 6 lets the tax team assemble a complete response in roughly 3 hours, drawing on the immutable audit trail, the TP knowledge graph's arm's length analysis, and the SHAP attribution reports. The same response would have taken 3–4 weeks manually.
Key Insight	The case illustrates ITIF's most distinctive capability relative to existing tax technology: continuous, AI-driven tax intelligence across a multi-system environment during an ERP transition. Standalone tax applications lose coverage when the legacy ERP has been partially cut over and the target ERP is not yet fully configured. ITIF's schema adaptation layer closes that gap by normalizing data from all systems simultaneously regardless of migration status.

A. The Parallel-Run Pattern for Tax Systems Migration

The parallel-run deployment approach is not specific to ITIF. It is a generalized pattern in which a new system operates in shadow mode alongside the legacy system, with tax teams comparing outputs and progressively increasing their reliance on the new system. ITIF's unified tax fabric suits this pattern well because it is built from inception to consume data from multiple heterogeneous sources at once. The Operational risk during parallel-run periods arises from the divergence between the determinations made by the two systems. Layer 2 confidence scoring and Layer 3 anomaly detection identify and escalate those divergences automatically rather than surfacing them during close. That automatic divergence detection is the single most operationally significant advantage over manual parallel-run approaches.

B. Why Existing Tax Technology Falls Short in the Same Scenario

A Level 2 organization running ONESOURCE or Vertex against SAP ECC in this scenario would lose tax determination coverage for Dynamics AX transactions until the new ERP's integration is certified—a process that typically lags the business cutover by several months. TP monitoring, still performed on periodic SAP extracts, would not detect the 847 mispriced APAC service-fee transactions until the quarter-end consultant review, well past the window in which remediation is inexpensive. Provision close would remain dependent on manual reconciliation across the two ledgers, and audit responses would revert to multi-week document assembly. The case therefore illustrates what a mature ITIF deployment enables that incremental point-solution investment does not: continuous, explainable tax intelligence through a period when source systems themselves are in flux.

IX. ARCHITECTURE EVALUATION

This section compares ITIF with two alternative tax technology paradigms: legacy tax systems (standalone applications with batch ERP integration) and rule-based tax engines (standalone applications with rule-based automation but no ML). The evaluation spans nine capability dimensions. Table III summarizes the comparison; Fig. 3 provides a visual representation.

TABLE III. TAX ARCHITECTURE COMPARISON: LEGACY vs. RULE-BASED vs. ITIF

Evaluation Dimension	Legacy Tax Systems	Rule-Based Engines	ITIF (This Paper)
TP Anomaly Detection	Manual review only	Threshold rules	Ensemble ML + SHAP
Real-time Tax Determination	No	Batch / near real-time	Yes
Regulatory Auto-Update	No	Partial	Yes
Cross-Jurisdiction Consistency	Spreadsheet-based	Rule engine	Jurisdiction knowledge graph
XAI / Explainability	No	No	Yes
Provision Automation (ASC 740)	No	Partial	Yes
RAG Audit Query	No	No	Yes
Schema Drift Resilience	No	Partial	Yes
ERP Migration Continuity	High risk of gap	Partial	Parallel-run support

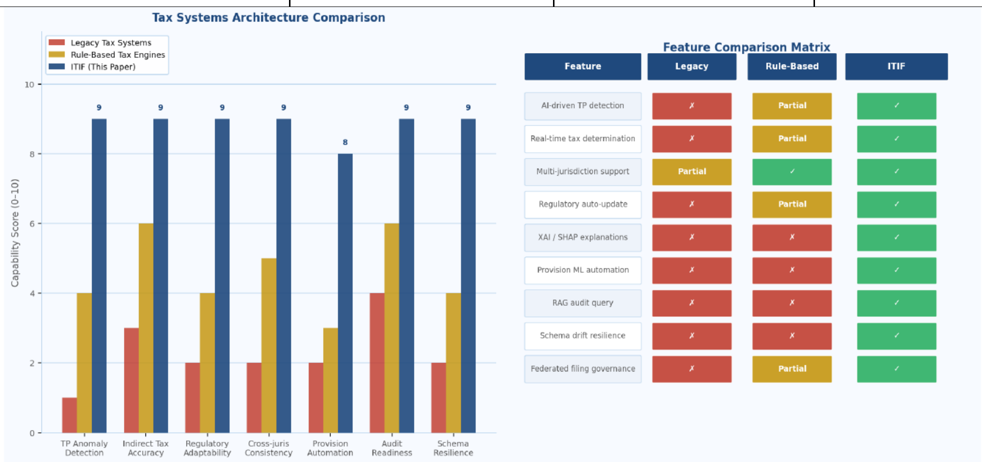

Fig. 3. Tax Systems Architecture Comparison—capability scores and feature matrix for Legacy Tax Systems, Rule-Based Tax Engines, and ITIF. Scores represent qualitative assessment, not empirical measurement.

A. Transfer Pricing and Provision Intelligence

The largest capability gaps in legacy and rule-based systems sit in the two hardest domains: transfer pricing and provision. Neither paradigm offers real-time TP monitoring, and neither supports ML-based UTP assessment; both fall back on periodic review and binary judgment. ITIF addresses both through the ensemble architecture of Layer 3 and the Monte Carlo capability of Layer 5. SHAP attribution deserves particular emphasis here. Tax authority examination requires that enterprises defend positions with specific facts; a TP flag without attribution is operationally inferior to one accompanied by documentation-ready SHAP output.

B. ERP Migration Continuity

Legacy and rule-based tax systems carry real continuity risk during ERP migration: their single-system batch integration lapses whenever the ERP is in transition. ITIF's multi-system data unification, enabled by schema adaptation across sources, is the only architecture evaluated that provides continuous coverage across that window. The capability is not cosmetic—it addresses one of the highest-risk periods in the enterprise tax calendar.

C. Regulatory Adaptability

The graph-based regulatory update mechanism, drawing on the knowledge graph construction methodology surveyed in [16], lets ITIF respond to regulatory change faster than rule-based systems. NLP-driven regulatory parsing can propose graph updates within hours of a tax authority announcement; manual rule updates in legacy systems typically take weeks or months. The advantage is most visible in indirect tax, where rate and exemption changes occur frequently and often in parallel across jurisdictions.

D. Evaluation Limitations

The capability scores here rest on architectural analysis and practitioner experience, not empirical benchmarking. The nine dimensions are themselves a testable hypothesis. Empirical validation of TP anomaly detection precision and recall, indirect tax determination accuracy against rule-engine baselines, and provision cycle time reduction would all strengthen the evidence base.

X. DISCUSSION

A. Positioning within Enterprise AI Research

ITIF sits at the intersection of two research programs that have evolved largely in parallel. Machine-learning-for-financial-compliance work [1, 16] has produced rigorous techniques for detection and mapping but has rarely been integrated into a coherent enterprise architecture. ML-systems architecture work [2, 14] has produced mature deployment patterns but has rarely been applied to specific regulated domains. ITIF treats tax as the domain in which those two programs have to converge. The framework's contribution is less the invention of new techniques than the disciplined integration of established ones into a production-shaped tax architecture.

B. The Role of External Tax Technology Vendors

ITIF is not a proposal to rebuild tax capability from scratch. Specialized vendors provide mature rule engines and data resources—rate tables, commodity libraries, TP benchmarking—that remain essential inputs. The framework's architectural contribution is to integrate those capabilities inside a coherent AI-native stack rather than replace them. Layer 4 uses Avalara or Vertex as the base rule engine and deploys the ML override only for the ambiguous long tail the rule engine handles poorly. That human-in-the-loop, ML-augmented-rule posture applies across all six layers and matters for practical adoption where vendor relationships are established stakeholders.

C. Limitations and Future Work

Three limitations warrant acknowledgment. First, ITIF is a conceptual framework without empirical validation; transfer pricing ensemble accuracy, indirect tax ML override uplift, and Monte Carlo UTP calibration all require controlled study. Second, the data unification layer assumes structured data availability; tax-relevant information often lives in contracts, invoices, and correspondence that demand additional NLP processing, potentially leveraging the Layer 6 RAG infrastructure as a document-understanding capability. Third, the framework addresses technical architecture while bracketing organizational and talent change. Adopting AI-driven TP monitoring and automated provision rewires how tax teams work, what skills they need, and how they engage external advisors—a transformation as significant as the technical one. Future empirical work on the TSMMM specifically should prioritize two programmes: (1) a multi-enterprise survey assessing TSMMM level placement across a representative sample of Fortune 1000 tax departments, using the binding-constraint criteria at each level as classification variables; and (2) a paired case study protocol contrasting one enterprise operating at TSMMM Level 2 and one at Level 4, providing depth to complement the survey's breadth. Both programs are feasible through existing practitioner research relationships in the tax technology community and are identified as priority empirical work for this series.

XI. CONCLUSION

ITIF offers a six-layer architecture for AI-driven enterprise tax systems modernization. It addresses data fragmentation, classification at scale, transfer pricing monitoring, indirect tax complexity, provision automation, and filing readiness through a coherent stack rather than a set of unconnected tools. Its most distinctive contributions are the transfer pricing intelligence module—an explainable, real-time ML monitoring capability calibrated to regulatory documentation standards—and the parallel-run continuity capability that closes the coverage gap during ERP migration.

The framework is positioned for practitioners and researchers alike. For practitioners, the TSMMM provides a staged path that matches investment to the binding constraint at each maturity level, avoiding the common failure mode of acquiring advanced capabilities before the data and governance foundations can support them. For researchers, the framework treats tax as a field where financial-compliance ML and enterprise ML-systems engineering must come together and points out the specific empirical questions—TP ensemble precision and recall, ML-override uplift in ambiguous indirect tax categories, and Monte Carlo UTP calibration—that need to be answered to significantly improve the evidence base.

The tax function of the next decade is not a team of spreadsheet specialists working after hours during the close. It is an intelligent, continuously operating, self-documenting compliance system. ITIF is a proposal for what that system's architecture should look like and an invitation to test, refine, and extend it against the operational realities of the enterprises that will operate it.

REFERENCES

[1] A. Mirani, "AI-Driven Compliance Intelligence: A Conceptual Framework for Enterprise Financial Governance," Journal of Computational Analysis and Applications (JoCAAA), vol. 35, no. 5, pp. 116–131, May 2026. [Online]. Available: https://eudoxuspress.com/index.php/pub/article/view/5467

[2] A. Mirani, "Designing AI-Native Financial Systems: Architecture Patterns for Intelligent Enterprise Platforms," IPHO-Journal of Advance Research in Science and Engineering, vol. 4, no. 4, pp. 21–33, 2026, doi: 10.5281/zenodo.20215472.

[3] A. Paleyes, R.-G. Urma, and N. D. Lawrence, "Challenges in deploying machine learning: A survey of case studies," ACM Computing Surveys, vol. 55, no. 6, Art. 114, pp. 1–29, Dec. 2022. [Online]. Available: https://doi.org/10.1145/3533378

[4] E. Ash, M. Guillot, and L. Han, "Machine extraction of tax laws from legislative texts," in Proc. Natural Legal Language Processing Workshop 2021, Punta Cana, Dominican Republic, Nov. 2021, pp. 76–85. [Online]. Available: https://aclanthology.org/2021.nllp-1.7/

[5] M. Weber, G. Domeniconi, J. Chen, D. K. I. Weidele, C. Bellei, T. Robinson, and C. E. Leiserson, "Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics," in Proc. KDD Workshop on Anomaly Detection in Finance, 2019, arXiv:1908.02591. [Online]. Available: https://arxiv.org/abs/1908.02591

[6] I. Chalkidis, M. Fergadiotis, P. Malakasiotis, N. Aletras, and I. Androutsopoulos, "LEGAL-BERT: The Muppets straight out of Law School," in Findings of the Association for Computational Linguistics: EMNLP 2020, pp. 2898–2904. [Online]. Available: https://aclanthology.org/2020.findings-emnlp.261/

[7] OECD, OECD Transfer Pricing Guidelines for Multinational Enterprises and Tax Administrations 2022. Paris: OECD Publishing, Jan. 2022. [Online]. Available: https://doi.org/10.1787/0e655865-en

[8] C. Coşkun, B. Gezici, M. Aydos, A. Tarhan, and V. Garousi, "ERP failure: A systematic mapping of the literature," Data & Knowledge Engineering, vol. 142, Art. 102090, Nov. 2022. [Online]. Available: https://doi.org/10.1016/j.datak.2022.102090

[9] Financial Accounting Standards Board, "ASC 740: Income Taxes," FASB Accounting Standards Codification, Norwalk, CT, 2023. [Online]. Available: https://asc.fasb.org/740

[10] International Accounting Standards Board, "IAS 12: Income Taxes," IFRS Foundation, London, 2023. [Online]. Available: https://www.ifrs.org/issued-standards/list-of-standards/ias-12-income-taxes/

[11] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Advances in Neural Information Processing Systems 30 (NeurIPS 2017), pp. 4765–4774. [Online]. Available: https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html

[12] A. Barredo Arrieta et al., "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," Information Fusion, vol. 58, pp. 82–115, Jun. 2020. [Online]. Available: https://doi.org/10.1016/j.inffus.2019.12.012

[13] P. Lewis et al., "Retrieval-augmented generation for knowledge-intensive NLP tasks," in Advances in Neural Information Processing Systems 33 (NeurIPS 2020), pp. 9459–9474. [Online]. Available: https://arxiv.org/abs/2005.11401

[14] Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency, "Supervisory Guidance on Model Risk Management," SR 11-7 / OCC 2011-12, Apr. 2011. [Online]. Available: https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm

[15] OECD, "Base Erosion and Profit Shifting (BEPS) Project," OECD, Paris, 2015–2023. [Online]. Available: https://www.oecd.org/tax/beps/

[16] A. Hogan et al., "Knowledge graphs," ACM Computing Surveys, vol. 54, no. 4, Art. 71, pp. 1–37, Jul. 2021. [Online]. Available: https://doi.org/10.1145/3447772

[17] M. Olbert and C. Spengel, "International taxation in the digital economy: Challenge accepted?," World Tax Journal, vol. 11, no. 1, pp. 3–46, 2019.

[18] S. Boyar, J. Roehl-Anderson, and D. Nellen, "Technology and the Tax Function," *Tax Adviser*, AICPA, vol. 52, no. 1, pp. 18–24, Jan. 2021.

[19] M. Kleppmann, Designing Data-Intensive Applications. Sebastopol, CA: O'Reilly Media, 2017. [Online]. Available: https://dataintensive.net/

[20] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, Design Patterns: Elements of Reusable Object-Oriented Software. Reading, MA: Addison-Wesley, 1994. [Online]. Available: https://isbnsearch.org/isbn/9780201633610

[21] Deloitte, "Tax Transformation Trends 2023: Operating Model Reimagined," Deloitte Tax LLP, 2023. [Online]. Available: https://www.deloitte.com/global/en/services/tax/research/tax-transformation-trends.html