

AI & Tax: we have got a story to tell

How can tax departments adapt to AI tomorrow?

March 18, 2026



AI & Tax, we have got a story to tell

This report is not a capabilities deck or a vendor pitch. It is an account of what our team has actually been doing: running simulations, building tools, forming and revising hypotheses, and positioning ourselves as clients to test our own advice. Some of our early assumptions did not hold. We have included those revisions, because we believe the thinking is more useful to you than the polish.

Artificial intelligence is transforming tax, yet most mid-market tax teams face a gap between vendor promises and operational reality. This report bridges that gap: it gives internationally active companies a grounded understanding of what AI means for tax operations today, what it does not yet solve, and where to start. It is a practical playbook that any tax team can begin implementing within weeks.

Before asking “which AI tool should we buy?” the more productive question is: “Do we understand the landscape well enough to make that decision wisely?” What is the difference between automation and AI, and where does each apply? Where are the genuine bottlenecks in our tax operations, and what governance is non-negotiable? A generative AI tool can draft a transfer pricing narrative in minutes, but it still needs the correct tested party, approved policy wording, and figures that reconcile to the general ledger, the provision, and Country-by-Country Reporting. Without a governed data layer and a deterministic calculation engine underneath, the AI is generating plausible-sounding text on an unreliable foundation.

We hope you will enjoy our insights into this interesting journey, and we will be sure to share many more going forward.

Best regards,

Mario van den Broek

Partner International Tax & Business Consulting Services RSM Netherlands

An overview of this report

This report is structured to move from diagnosis to build plan. It is organized in five chapters and uses a five-layer “Modern Tax Department” model as the lens to identify bottlenecks and define minimum deliverables that can be implemented in controlled steps.

The model is read bottom-up: (1) infrastructure and controls, (2) data and documentation, (3) standard processes and deterministic calculators, (4) automation and AI execution, and (5) outputs and obligations. Each layer is translated into practical deliverables and control points—so progress can be made incrementally while maintaining auditability, segregation of duties, and a clear “single source of truth” for numbers and narrative.

Chapter 1 defines what is changing in tax delivery and what “good” looks like operationally. Chapter 2 clarifies automation versus AI and how they work together in tax workflows: deterministic logic produces the numbers, automation executes repeatable steps and checks, AI supports drafting and detection, and human review remains the final control point.

Chapter 3 provides the empirical backbone through workflow simulations and scenario walkthroughs. It shows where time is typically lost (data chasing, reconciliation loops, reruns, and documentation rebuild), why rework occurs, and which controls reduce operational risk without slowing the cycle. Chapter 4 turns the model into action: what a first pilot looks like end-to-end, which minimum deliverables must exist to run it safely, and how to scale only what remains stable, explainable, and rebuildable.



1

Why Tax and AI matter now

Rising complexity and limited team capacity



2

Automation versus AI

Explanation and working together



3

Simulation Revelations

Where time is lost today and where gains appear



4

How to start tomorrow

A simple path to begin and scale safely

Intended audience



Revenue

€/\$100M–3B




Footprint

5–40 countries, multiple legal entities



HQ

US, Europe, or Asia



Tax Team

Central oversight with local support


This report is intended for central tax functions within mid-market internationally active companies. The remit is broad: corporate income tax, transfer pricing, VAT, and global employer services, sometimes shared across a team of five to fifteen.

Specific roles: Head of Tax or Group Tax Director (strategic priority-setting and budget justification), Transfer Pricing Manager (operational workflow and monitoring), Tax Technology Lead (infrastructure and data layer), CFO with Tax Oversight (investment case and cross-functional alignment).

The report focuses on tax-led improvements that deliver results within existing budgets—without requiring large technology programmes, dedicated AI budgets, or teams of data scientists. The operating model described here is designed for teams that want progress, not perfection, and that measure success by whether the next close cycle is faster, cleaner, and more defensible than the last one.

The report is written for teams that want progress without launching a large technology program. It focuses on tax-led improvements that can be built with existing tooling and controlled steps. The objective is repeatable and traceable delivery, so time moves away from rework and towards exceptions, judgement, and decisions that materially affect tax outcomes.

Understanding our working hypothesis



Our approach reads from bottom to top whereby each layer enables the one above.

Step 1. Build the foundation

Infrastructure and controls set access, logging, security, and audit readiness.

Step 2. Create shared memory

Data and documents become one trusted source for numbers and evidence.

Step 3. Standardize workflows

Tax processes and calculators turn shared inputs into consistent outcomes.

Step 4. Add intelligent execution

Automation and AI run steps, draft content, and route exceptions to people.

Step 5. Deliver defensible outputs

Business and obligations produce filings, reports, and responses with traceability.

Bottom line

Bottom layers control quality. Top layers deliver outcomes.

Understanding our working hypothesis

Layer 1: Infrastructure and controls

This is the technical foundation. It makes everything above safe, compliant, and maintainable. It covers connectors to source systems, access management, security, logging, monitoring, and cost control. Without this layer, outputs are hard to audit and hard to sign off.

Layer 2: Data and documents

This is the shared tax memory. It combines a structured tax data layer with an approved document corpus. It solves data and evidence being scattered across systems, inboxes, and drives. With this layer, everyone works from the same numbers and the same sources.

Layer 3: Tax processes and calculators

This layer standardizes workflows for CIT, VAT, TP, and payroll. It replaces spreadsheet logic with deterministic calculators and reusable templates. It reduces local variants, rework, and inconsistency. The same inputs produce the same outcomes.

Layer 4: Automation and AI layer

This layer runs intelligent execution on top of shared data and standard workflows. It handles routine steps such as checks, reconciliations, and first draft documentation. It uses orchestration, specialized agents, and a review queue. Exceptions go to people. Judgement stays with the team.

Layer 5: Business and obligations

This is what the tax team must deliver. It includes filings, reports, dashboards, and audit responses. This layer becomes defensible when it is fed by controlled data, standard workflows, and documented review. Outputs are complete, consistent, and ready for scrutiny.



Each layer removes a specific risk. Together they create a controlled production model for tax.

01

RSM

Why Tax and AI matter

Why AI & Tax is a “thing” now?


Tax is moving from a periodic reporting function to a continuous, data-tested control function. With electronic invoicing and Standard Audit File for Tax (SAF-T) style data requests, authorities increasingly validate filings against transaction-level data they can obtain quickly. That changes the timing of scrutiny: questions arrive during the period, not months later, and they expect you to rebuild the number from source to filing without a “spreadsheet story” in between.

OECD Pillar Two adds a second pressure point. It turns tax into a recurring data exercise across entities, requiring consistent definitions, stable mappings, and reconciliations that can be rerun when data updates. If each cycle still starts with chasing exports, aligning entity identifiers, and reworking calculations because formats differ, the team’s capacity gets consumed by rebuilding the same foundation every time. The result is predictable: late close stress, inconsistent outcomes between provision, return, and management reporting, and increased reliance on individual workarounds.


Stakeholder expectations have also shifted. Boards and chief financial officers want earlier visibility into effective tax rate drivers, cash tax movements, and exposures so decisions can be made before quarter-end. Auditors want traceability: what data was used, which logic produced the numbers, what changed since last cycle, and who approved exceptions. When evidence is scattered across inboxes and versions drift across files, defensibility becomes the bottleneck.

This is why automation and artificial intelligence matter now. Automation makes the process repeatable by pulling data, applying the same mappings, running deterministic calculations, and flagging failed tie-outs early. Artificial intelligence supports the “explain” layer by drafting narratives from approved inputs and summarizing exceptions for review. The outcome is not more technology for its own sake; it is a tax process that can withstand real-time scrutiny without heroics.

A note on our own starting point: when we began this work, we assumed clients would be excited, that the bottleneck would be technical, and that adoption would follow naturally once the tools worked. Almost none of this held. The surprises were human, not technical. We have built this report around what we actually found, not what we expected to find.



The tax reality landscape



A shift to real-time, data-driven tax supervision

- Mandatory digital reporting (SAF-T, e-invoicing/CTC, Pillar Two, DAC6–8) is accelerating.
- Authorities expect structured data, automated validations, and end-to-end traceability.

Fragmented systems across the organisation

- Transaction data scattered across ERPs, billing, payroll and local tools.
- Supporting evidence spread across e-mails and shared drives.
- No single environment where data, logic and documentation converge.

Rising audit expectations

- Every reported number must be reproducible and tied back to a clear rule and source.
- Provision, return and TP documentation must be fully aligned.
- Variances and exceptions must be explainable within minutes, not days.

The impact

A widening gap between **what authorities expect** and **what current manual workflows can reliably deliver**.

Client adoption and demand

Adoption trends:

- Over 58% of corporate tax departments report feeling under-resourced and are actively exploring automation and AI solutions.
- Around 45% of large enterprises have started AI pilots in data reconciliation, VAT validation, or document review.
- Generative AI exploration is accelerating, 67% of firms have yet to deploy, but more than half plan inclusion in their 2025–2027 roadmap.
- Across EMEA, AI is most used for Pillar Two data processing, e-invoicing, and transfer pricing document generation.


Client priorities:

- Efficiency:** Reducing compliance workload and manual reconciliation (30–40% time savings in early adopters).
- Accuracy & Standardization:** AI-driven validation improves consistency across multi-jurisdictional filings.
- Regulatory Readiness:** Firms are preparing for the EU AI Act's transparency and documentation mandates.
- Scalability:** Demand for scalable AI tools that integrate across finance, ERP, and reporting systems.

RFP & market demand signals:

- Growth in AI-related RFPs for tax compliance software and data transformation projects (+25% YoY in 2024).
- Clients increasingly request AI explainability and audit trail features as standard procurement criteria.
- Notably, cross-functional RFPs (tax + finance + IT) are becoming more common, signaling organizational alignment in AI adoption.

AI Adoption in Tax Departments (2024-2025)



Source: Thomson Reuters State of the Corporate Tax Department , 2025).

02

RSM

Automation vs AI



Automation in more detail

Automation is rule-based execution. If you can write the steps down, software can run them the same way every cycle. It is ideal for structured work such as data pulls, entity and account mapping, standard calculations, and validation checks. Robotic process automation (RPA) is a common example. Automation is fast, consistent, and traceable, but it does not handle judgement or ambiguous exceptions.


AUTOMATION: THE DIGITAL WORKHORSE

Rule-Based Execution of Repetitive Tasks

ROBOTIC PROCESS AUTOMATION (RPA)




Software bots performing tasks exactly as instructed without human intervention. It doesn't "think," it follows predetermined rules.



Example: Automatically copying-pasting figures, completing forms in seconds.

TYPES & MECHANISMS

DETERMINISTIC LOGIC ENGINES



Uses fixed formulas (e.g., deferred tax calculation) or predefined rule sets.

AUTOMATED VALIDATION CHECKLISTS



Automatically verifies required fields are filled or data formats are correct.

KEY CHARACTERISTICS & APPLICATION

RELIABLE & PREDICTABLE



FAST & CONSISTENT: Excels at structured, highly repetitive tasks.



DIGITAL WORKHORSE: Automates data entry, workflow routing, and status updates with 100% consistency.



DOES NOT THINK: Cannot handle exceptions outside explicit programming.



IN SHORT: If you can write down the step-by-step rules, you can automate it. (e.g., Excel Macros, Data Scripts).

AI in more detail

AI refers to software that exhibits a degree of “intelligence” – it can make inferences or decisions that aren’t explicitly hard-coded, often by learning from data or by mimicking human reasoning. A common branch of AI is machine learning (ML), where an algorithm finds patterns in historical data and uses them to flag anomalies or make predictions.


In a tax context, an ML-based tool might analyze thousands of transactions and learn what “normal” looks like, then flag exceptions in large datasets that deviate from the norm. For example, it could scan all invoice postings and alert you to ones that look unusual (maybe a purchase logged with an uncommon tax code or an amount far outside expected ranges) without you pre-defining every scenario.

Another branch of AI includes large language models (LLMs) or generative AI, which can draft and summarize content. In tax, an LLM-based assistant might generate a first draft of a tax technical memo or a portion of a tax provision narrative, using context you provide. It won’t be 100% correct or final, but it gives the team a head start on documentation.

AI is thus useful for less structured tasks, those that involve recognizing patterns, making judgment-like calls on data, or creating narrative text. It brings a degree of flexibility and “smarts” that pure automation lacks. However, AI’s outputs are probabilistic, meaning they require review and oversight; the AI might flag 100 transactions and 5 of those could be false alarms, or it might draft a memo that needs editing for accuracy.

ARTIFICIAL INTELLIGENCE (AI) IN TAX: LEARNING, INFERENCE & FLEXIBILITY

1. AI: BEYOND HARD-CODED RULES



Software that exhibits “intelligence,” making inferences or decisions by learning from data or mimicking reasoning.



MACHINE LEARNING (ML): PATTERNS & ANOMALIES.
Algorithm finds historical data patterns.



TAX EXAMPLE: Scans thousands of transactions to flag **UNUSUAL POSTINGS** (e.g., uncommon tax codes, unexpected amounts) without pre-defined rules.

2. GENERATIVE AI (LLMs): CONTENT & DRAFTING



Large language models that draft and summarize content.



TAX EXAMPLE: Generates a first draft of a **TAX TECHNICAL MEMO** or provision narrative, providing a head start.

3. KEY CHARACTERISTICS & ROLE: FLEXIBLE & PROBABILISTIC



USEFUL FOR UNSTRUCTURED TASKS.
Handles pattern recognition, judgment-like calls, and narrative creation.



PROBABILISTIC OUTPUTS.
Requires review; might flag false alarms or draft with inaccuracies.



HUMAN OVERSIGHT ESSENTIAL.
AI brings flexibility & “smarts” but demands human final review for accuracy.

Automation and AI hand in hand

In practice, automation and AI work best when they are combined in one controlled workflow. Automation provides consistency and hard controls. A deterministic tax calculator applies the same logic for every entity. An RPA bot pulls the same inputs, the same way, every time. AI adds value where rules start to struggle. A machine learning model can flag unusual postings that a simple rule might miss. A writing assistant can draft the first version of tax footnotes or an ETR bridge narrative based on the numbers.

The strongest setup is simple: **numbers from automation, language from AI**. The workflow stays auditable because AI is only one step inside a broader process with clear inputs, outputs, and review points. A practical example is an agent that pulls trial balance data (automation), runs standardized tax formulas and validations (automation), and then asks an AI assistant to draft an explanation of the outcome (AI). A human reviews and signs off before anything is used externally.

The takeaway is that automation handles repeatable procedures, AI supports analysis and drafting, and human judgment remains the final control gate. This distinction helps teams decide when a straightforward script is enough and when AI can add value without weakening control.

A central illustration shows two hands shaking. The hand on the left is a mechanical arm with a silver and blue finish, labeled "AUTOMATION" on its forearm. The hand on the right is a glowing, translucent blue digital hand with circuit-like patterns, labeled "AI" on its forearm. They are set against a dark blue background with glowing circuit lines and icons like a document, a bar chart, and a pie chart.

WORKING TOGETHER

HUMAN OVERSIGHT

03



Practical observations through simulation

The technology mostly did what it claimed to do. The surprises were human.

Corporate tax simulations

Provision and effective tax rate work often starts with a local Excel model and ends with a consolidation exercise that lands late in the cycle. Inputs arrive in waves, often after the close is already moving, so group tax spends time normalizing formats, rekeying adjustments, and reconciling versions instead of analyzing outcomes. Under time pressure, the effective tax rate bridge becomes a manual reconstruction, so movements are explained as a narrative rather than a traceable story linked to defined drivers. The practical result is a longer close and more last-minute questions from the Chief Financial Officer and auditors, because the numbers can be produced but not consistently rebuilt.

Compliance and returns are typically managed country by country with locally grown workbooks, checklists, and filing packs. That works until the group needs consistency across jurisdictions, for example when aligning provision positions with what is ultimately filed, or when ensuring incentives and disclosures are handled the same way in each pack. Headquarters then becomes the clearing house for missing information, follow ups, and manual tie outs between “provision logic” and “return logic.” The time cost is less the technical tax work and more the repeated coordination and reconciliation required to get everyone onto the same definitions, evidence, and final numbers.

Uncertain tax positions, audits, and disputes are often tracked across separate spreadsheets, email threads, and memos, each with its own status and assumptions. This makes it difficult to maintain one current view of what is open, what changed since last quarter, and what the likely financial statement impact is at group level. When questions come in from management or auditors, the team rebuilds the picture by collecting updates from countries and checking which version is current. The gap is not that the work is not done, but that there is no live register that links exposures, timelines, decisions, and supporting evidence in one place.

Forecasting and planning frequently treats tax as an output of the financial forecast rather than a modelled component of it. In practice, tax is applied as a flat percentage, which hides the drivers that move cash tax and the effective tax rate over time. Expiring losses, changes in profit mix, restructurings, and new regimes such as Pillar Two are hard to stress test if the forecast does not reuse the same data definitions and calculation logic as provision and returns. The consequence is lower confidence in forward looking effective tax rate and cash tax, and tax re enters the conversation late when the business needs answers early.



Corporate tax simulations

We walked through a typical corporate income tax (CIT) cycle using the artefacts most tax functions already have: an enterprise resource planning (ERP) trial balance export, a tax adjustment schedule, the prior year provision workbook, the prior year return pack, and the current list of uncertain tax positions (UTPs) and open audits. We mapped where data is re keyed, where reconciliations happen, and where reviews and sign offs sit today. We then reran the same cycle assuming one shared mapping, one standard calculator, a small set of automated checks, and a simple run log with approvals.

What our simulations reveal in practice. The bottleneck is rarely technical tax work. It is input quality, reconciliation, and rebuilding evidence after the fact. In a typical 15-entity group, the current CIT close requires approximately 12–18 person-days of reconciliation per quarter. Under the controlled model with shared data and standard calculators, simulation shows this reduces to 4–6 person-days—a 60–70% reduction in manual effort. The key driver is elimination of version drift between provision and return.

CIT after automation and artificial intelligence (AI): The defining shift for corporate tax is the elimination of version drift between provision and return. Unlike other domains, CIT depends on a single reconciled data chain from general ledger to filed return—today that chain breaks at every handoff. The controlled model rebuilds it as one continuous flow:

- **Shared tax data layer plus standard calculators:** GL and adjustment data for all entities feed a single CIT provision engine. Book-to-tax, current and deferred tax, and the ETR all derive from one source—eliminating the 3–5 days per quarter typically spent reconciling provision to return.
- **Standardized CIT compliance pack per entity:** Each entity completes the same digital pack for trial balance, adjustments, calculations, and return mapping. Validations and checklists enforce quality. The same numbers feed returns, Country by Country Reporting (CbCR), Pillar Two global minimum tax reporting, and management reporting. This reduces chasing and reconciliation effort.
- **Integrated UTP and controversy dashboard:** All uncertain positions and audits sit in one register linked to provision and returns. CFOs see real-time exposure by entity and jurisdiction—reducing quarter-end UTP assessment from 2–3 days to a continuous, always-current view.
- **CIT forecasting and scenarios on the same platform:** Provision and compliance logic is reused on forecast data. Tax can model future ETR and cash tax by country, test rate changes and Pillar Two scenarios. Forecast generation reduces from 2–4 days to hours when the model already exists.


Transfer pricing simulations

Year-end fire drills and unplanned true-ups. In many groups, transfer pricing performance is only reviewed when the year is already closing. Margins are checked late, using data that arrives in waves and is not always consistent across countries. The consequence is that required adjustments appear suddenly, bookings need to be made under time pressure, and the story behind the numbers is reconstructed after the fact. This is exactly when the Chief Financial Officer and auditors ask for a clear explanation, and the team ends up managing surprises instead of running a controlled cycle.

Fragmented policies and data. Local teams often maintain their own profit and loss data, spreadsheets, and benchmarking files. Over time, that creates “local variants” in method selection, tested party choices, arm’s-length ranges, and allocation keys. Even when everyone intends to follow the same policy, differences in definitions, inputs, and versions make outcomes hard to compare and hard to defend consistently. The work becomes less about transfer pricing judgement and more about reconciling who used which file, which assumptions, and which numbers.

Heavy documentation workload. Local File deliverables, Master File sections, and intercompany agreements are frequently rebuilt each year from scattered sources. Narrative text is copied, updated, and revalidated across documents, while supporting evidence sits in different folders and email threads. Because the documentation is not structurally linked to the general ledger and the same underlying datasets, numbers do not always reconcile cleanly to the provision, Country-by-Country Reporting, or management reporting. The result is repeatable rework, plus a higher risk that inconsistencies show up during review.

Limited in-year insight. Transfer pricing then becomes mainly a compliance exercise instead of a management tool. There is limited ability to monitor margins during the year, predict exposures early, or support business decisions proactively (for example pricing changes, restructurings, or supply chain shifts). When issues are only visible at year-end, the organization is forced into reactive corrections, rather than controlled adjustments made early with a clear rationale and audit trail.




Transfer pricing simulations

What it is based on. We walked through a typical year-end transfer pricing cycle using artefacts most groups already have: periodic entity profit and loss extracts, intercompany transaction detail, the group transfer pricing policy and benchmark files, the prior-year true-up workbook and journal entry support, and prior Local File and Master File content. We mapped where margins are only assessed at year-end, where country teams apply different methods or ranges, where documentation is rebuilt from scattered sources, and where numbers do not tie cleanly to finance outputs. We then re-ran the same cycle assuming one shared dataset, one controlled policy file, continuous monitoring, and review gates with a simple run log.

What our simulations reveal in practice. The bottleneck is rarely the transfer pricing judgement itself. The effort sits in collecting late data, reconciling versions, and reconstructing documentation year-end. In a group with 20+ intercompany flows, the annual TP cycle typically absorbs 25–40 person-days. Under continuous monitoring with standardized data feeds, simulation shows this reduces to 8–15 person-days—a 55–65% reduction.

Transfer pricing after automation and AI: The defining shift is continuous margin monitoring replacing the year-end fire drill. Unlike CIT (version integrity) or VAT (upstream coding), TP's transformation turns a once-a-year compliance exercise into an in-year management tool:

- **Continuous margin monitoring with alerts:** A central TP monitor pulls periodic P&L data per entity, compares actual margins to policy ranges, and flags deviations before year-end. This eliminates the typical 5–8 day scramble to assess whether true-ups are needed.
- **Policy as code and consistent true-ups:** The transfer pricing method, arm's-length range, and allocation keys are encoded once and reused. The same deterministic rules calculate margins, determine the true-up required, and generate booking instructions with a traceable audit trail of inputs, versions, and approvals.
- **Standardized transfer pricing data and documentation pack:** One structured dataset feeds Local File and Master File schedules. AI drafts functional analyses using approved templates, reducing documentation time from 3–5 days per entity to under 1 day. Human review remains the final gate.
- **In-year transfer pricing cockpit for finance and tax:** A single dashboard shows current margin positions, expected year-end outcomes, and actions needed. This turns transfer pricing from backward-looking compliance into a forward-looking management instrument.



Value Added Tax (VAT) simulations

VAT Compliance Challenges



Multi-country returns and reconciliations. Value-added tax compliance is often run country-by-country, using local spreadsheets or local advisors with their own data requests and formats. Head office then spends time chasing inputs, resolving late differences, and tying value-added tax control accounts in the general ledger back to what was filed locally. The cycle becomes a coordination and reconciliation exercise, not a controlled production process.

Weak controls and Accounts Payable miscodings. Period-end checks are frequently sample-based and manual, which means errors survive into filed returns. Mis-posted tax codes, missed reverse charges, and non-deductible value-added tax are often only discovered during an audit, or after filing when corrections are costly and explanations are time-consuming. The practical impact is rework, uncertainty in the numbers, and a recurring fear of “what will the tax authority find first”.

Cross-border complexity and extra reporting. Place-of-supply rules, reverse charge mechanics, and cross-border flows create a high volume of edge cases that are hard to manage with ad-hoc workbooks. EC Sales Lists, Intrastat, and digital reporting requirements such as e-invoicing and Standard Audit File for Tax are frequently managed as separate streams, with different datasets and owners. As a result, inconsistencies and omissions are common, and the team rebuilds evidence across multiple files.

Registrations and system set-up. New countries, rate changes, and changing reporting mandates often require manual enterprise resource planning updates and local workarounds that are not centrally controlled. Registration thresholds and new obligations are monitored intermittently, so risks are identified late and remediation becomes reactive. In practice, value-added tax set-up and compliance readiness lag business growth, even when the underlying data exists.

Value Added Tax (VAT) simulations

What it is based on. For the simulation, we reconstructed a multi-country value-added tax (VAT) month end using the materials most teams already have: enterprise resource planning (ERP) transaction extracts for sales and purchases, Accounts Payable (AP) tax code master data, the VAT control account reconciliation to the general ledger (GL), and last period's return pack per country. We also included the parallel outputs that must line up, such as EC Sales Lists, Intrastat declarations, and digital reporting files (e-invoicing and Standard Audit File for Tax). We mapped where teams chase missing inputs, correct miscoding, and redo tie outs, then re-ran the same cycle assuming one governed dataset, centrally maintained rules, automated validations, and a simple run log with approvals.

What our simulations reveal in practice. VAT performance is decided upstream. When coding is inconsistent and checks are only run at filing time, errors are embedded before anyone reviews them. In a 10-country operation processing 50,000+ monthly transactions, manual VAT reconciliation typically takes 8–12 person-days per month. With upstream coding control and automated validation, simulation shows this reduces to 2–4 person-days—a 65–75% reduction.

VAT after automation and AI: The defining shift is upstream coding control. Unlike CIT (version integrity) or TP (continuous monitoring), VAT's transformation happens before the return—at the point where transactions are coded in AP and the ERP. Fix the input, and the output fixes itself:

- **Global VAT compliance hub:** One platform ingests transactions from all systems, applies country rules, and prepares each return from the same data—replacing country-by-country preparation that typically adds 1–2 days per jurisdiction per month.
- **Continuous VAT control in Accounts Payable:** Automated checks validate tax codes, reverse charge treatment, and non-deductible VAT across the full population. AI highlights anomalies before month-end—catching errors that sample-based manual checks miss in 40–60% of cases.
- **Cross-border engine and aligned digital reporting:** A central place-of-supply engine determines correct VAT treatment for cross-border flows. The same dataset populates e-invoicing, SAF-T, and filing data—ensuring consistency across reporting streams that regulators increasingly cross-reference.
- **Registrations and change monitoring built in.** A registration monitor tracks thresholds and new market activity. Rate changes and rule updates are maintained centrally and applied consistently, which reduces local workarounds and late surprises.

Global Employer Services (GES) simulations

International hires and moves

Onboarding a cross-border hire is rarely one workflow. Human Resources, immigration, payroll, and tax each run their own steps and their own checklists. The dependencies are tracked in email threads and spreadsheets, so it is easy to miss a registration, a payroll cut-off, or a required tax identifier. The result is predictable: people start work while the back office is still catching up, and the first month becomes a correction exercise instead of a clean start.

Global payroll cycles

Each country runs payroll on a different calendar, often on different systems, with manual uploads from Human Resources and Excel. Late inputs and late changes drive net pay errors, off-cycle corrections, and rushed wage tax filings. When something goes wrong, the team spends time reconstructing what changed and why, rather than working the exception once and closing it. This is where confidence erodes, because the same issues reappear every month in a slightly different form.

Remote and flexible work

Remote work and frequent travel create compliance exposures that are hard to see without reliable location data. Thresholds, treaty positions, and social security treatment are often checked only after the fact, because the team is not monitoring patterns in-period. That is how shadow payrolls and host-country obligations are discovered late, sometimes only after an authority notice or an internal audit request.

Limited visibility and control

Leadership often lacks one view of who is working where, which payrolls have been processed, what has been filed, and where the compliance risk sits. Mobility and payroll specialists therefore spend their capacity on firefighting rather than improving employee experience and advising the business. Costs and risks become surprises, not managed outcomes, because the operating model does not produce a traceable end-to-end record of decisions, inputs, and approvals.

Global Employer Services (GES) simulations


What this is based on. We traced a full cross-border hire and pay cycle, from onboarding to first payroll close and first filings. We used the artifacts most teams already have: Human Resources Information System (HRIS) master data, payroll calendars and input files, immigration case trackers, travel or expense data as a proxy for work location, and prior wage tax and social security filings. We mapped where information is rekeyed, where approvals sit, and where issues typically surface late (missed registrations, shadow payroll, off-cycle corrections). We then re-ran the same cycle assuming one shared data layer, a standard playbook, automated validations, and a simple run log with sign-offs.

What our simulations reveal in practice. The recurring failure point is not the complexity of the rules. It is the fragmented data: worker record discrepancies, work location gaps, payroll entity mismatches, and filing status errors. A typical cross-border hire currently requires 15–20 person-hours of coordination across HR, immigration, payroll, and tax. Under a unified people-tax data layer, simulation shows this reduces to 5–8 person-hours—a 55–65% reduction.

Global employer services and payroll after automation and AI: The defining shift is the location-based compliance radar. Unlike CIT, TP, or VAT—which process known transaction flows—GES must track where people are, which changes constantly. The transformation centers on turning fragmented HR-payroll-tax data into one real-time view:

- **One shared people tax data layer with controlled documentation.** Human Resources, payroll, immigration, and travel or location inputs feed one worker record per person, linked to contracts, rulings, and policy evidence. Standard validations check completeness (start date, work location, tax identifier, payroll entity), and every change is logged for audit readiness.
- **Standard onboarding and payroll playbooks, executed end-to-end.** A single workflow orchestrates hires and moves from registrations through payroll input, checks, and approvals. Country differences are handled through parameters (for example wage tax setup, the 30% ruling, and social security certificates such as the A1 form), so the process stays consistent while outcomes remain locally correct.
- **Mobility and remote-work compliance radar: Work location signals are monitored in-period against tax and social security thresholds. The system alerts before threshold breaches—not after. This predictive capability distinguishes GES transformation: compliance risk moves with the person, so the monitoring must be continuous and location-aware.**
- **Governance and visibility by design.** Payroll close has explicit review gates, an exception register, and a run log that ties inputs to outputs. Leadership gets a live view of who is working where, what has been filed, and what risks are open, so compliance becomes controlled delivery instead of recurring firefighting

Start small, stay in control



The one rule

Do not scale what cannot be explained and rebuilt. Everything else follows from this.

Five moves. In this order.

- Define your data: one dictionary, one mapping, one owner
- Pick one pain point that repeats every close. One workflow. Not three.
- Run one end-to-end cycle: pull, validate, calculate, draft, review, sign off
- Build your tax memory in one place: version-controlled, read-only when final
- Scale only what holds: copy the pattern, never the chaos

Hard constraints

- Deterministic logic produces every number. Always.
- AI drafts and flags. It does not decide.
- A human signs off every output. No exceptions to the exception rule.

The hidden bottleneck

- Most pilots do not fail because the technology fails. They fail because no one decided who owns the output, what “correct” means, or when to override.
- Prompting is a professional skill. Output quality depends on question quality, not model quality. Build this into how your team works from day one.

Minimum deliverables



Tax Data Dictionary

Glossary of tax-related numbers



Tax-Owned Mapping Layer

Standardizes entity IDs and accounts



Tax Library with Version Control

Centralized repository for tax documents



Deterministic Calculator / Standard Pack

Locked template for repeatable outputs



Review Gates + Run Log

Control layer for tracking changes and approvals

1) Tax Data Dictionary

Start from the last provision pack, last return pack, and one trial balance export. Create one Excel tab with 25 key fields.

Columns: Field, Definition, Source report or tab, Validation tie out, Owner.

Rule: one owner approves definition changes.

2) Tax-Owned Mapping Layer

Start from the trial balance account list and the entity list used in consolidation.

Create two tabs: Entity map and Account map.

Account map focuses on tax-sensitive accounts first.

Columns: Local account, Description, Tax category, Treatment note, Owner.

Rule: any unmapped account is flagged and must be categorized.

3) Tax Library with Version Control

Start by collecting final filed returns, signed workpapers, rulings, and standard templates.

Create one SharePoint library with folders by tax type and year.

Keep drafts separate from finals.

Rule: finals are read-only and version history stays on.

4) Deterministic Calculator or Standard Pack

Start from the current best workbook used for ETR, VAT, or TP.

Rebuild as one locked template with tabs: Inputs, Calc, Checks, Outputs, Change log.

Move rates and assumptions to a Parameters tab.

Rule: any logic change creates a new version and needs a second review.

5) Review Gates plus Run Log

Start by writing the workflow steps on one page.

Gates: data tie out, calculation run, exceptions reviewed, output approved.

Create a run log row per cycle.

Fields: period, data timestamp, mapping version, pack version, overrides, reviewer, approver, output link.

Rule: no output is used externally without a completed log and sign-off.

How to start tomorrow

You do not need a budget. You need a decision.

What the research tells us

This report has been about evidence. We will end with conviction.

The tax profession is approaching a bifurcation. Those who have integrated AI into their judgment process and those who have not. The gap is invisible now. It will become visible in speed, depth, and the ability to see across jurisdictions and domains simultaneously—in every close cycle, from here forward.

The deepest bottleneck in this transition is not technical. It is epistemological. Tax professionals do not lack tools. They lack a new theory of what it means to know something professionally: how to form a view when the model is uncertain, how to trust output they did not produce, how to own advice that was partially generated. The governance artefacts in this report address the operational layer. The epistemic shift is the work underneath them—and it cannot be bought or installed.

The firms that win are not the ones with the best tools. They are the ones who help their clients think better in a world where thinking has become the scarce resource. Tools are table stakes. Judgment is the differentiator.

We have been building toward this moment for two years. The experiment is ongoing. The invitation is open. RSM Netherlands' International Tax Services team is ready to think alongside you.

Contact

Want to know more about the possible impact for your company?
Please do not hesitate to contact us



Mario van den Broek

Tax Partner
mvdbroek@rsmnl.nl



Mourad Seghir

Manager
mseghir@rsmnl.nl



Danillo Bujan Rodriguez


Consultant
dbujanrodriguez@rsmnl.nl

Please let us know what you think



THE POWER OF BEING UNDERSTOOD

ASSURANCE | TAX | CONSULTING



THE POWER OF BEING UNDERSTOOD
ASSURANCE | TAX | CONSULTING



RSM Netherlands Holding N.V.

Mercuriusplein 9
Postbus 30
2130 AA Hoofddorp
T 023 530 0400
rsmnl.com

RSM Netherlands Holding N.V. is een onafhankelijk lid van het RSM netwerk en handelt onder de naam RSM. RSM is de handelsnaam van elk bij het RSM netwerk aangesloten kantoor. Het RSM netwerk is een samenwerkingsverband van onafhankelijke accountants- en advieskantoren, die ieder voor eigen verantwoordelijkheid werken. Het RSM netwerk bestaat in geen enkel rechtsgebied als afzonderlijke rechtspersoon. Op alle opdrachten zijn onze algemene voorwaarden van toepassing, die op verzoek kosteloos worden toegezonden. RSM Netherlands Holding N.V. is ingeschreven bij de Kamer van Koophandel onder 34255762.

© RSM International Association, 2025