

Designing the future of work
A vision for generative AI strategy in
the tax and finance functions

April 2025

Introduction

Picture this: Eighteen months from now, the tax and finance functions of a large, globally dispersed organization are operating with a level of clarity, focus, and productivity that feels almost futuristic today. Regulatory submissions are turned around in hours, not days. **Tax technical memos are generated with embedded citations from internal policies and global legislation.** Budget commentary is auto-drafted from forecast data. Reconciliations run overnight. Audit evidence is auto-organized. Knowledge once locked in individual inboxes is accessible enterprise-wide, augmented by a private, secure artificial intelligence (AI) system.



This is not science fiction. This is what a successful generative artificial intelligence (genAI) implementation could look like—if done thoughtfully, strategically, and with the right vision. But unlike many one-off deployments that fade into the background, these genAI platforms are designed to evolve. **It is built to flex with the pace of change—because in a landscape where foundational models are improving every quarter, and new regulatory challenges emerge monthly, any AI solution that cannot adapt is a future liability.**

A strategy that starts with purpose, not procurement

Before a single model is tested or a single workflow built, the Chief Financial Officer (CFO) and Global Head of Tax convene a cross-functional working group. They begin not with a product demo but with questions:

- What business problems do we want AI to solve?
- What is our tolerance for risk and regulatory scrutiny?
- How will we keep this platform responsive to market and technological change?

These questions shape a genAI vision that balances ambition with realism. They agree that AI should not just reduce effort—it should expand capability, deepen insight, and future-proof the organization's particularly complex and costly functions.

A capability map that reflects the real work

The team launches a discovery sprint. Instead of asking what AI can do, they ask what the business needs it to do.

Over four weeks, they build a map of more than ninety use cases drawn from direct observation of work in tax, finance, audit, and treasury. These include both generative workflows and automation flows:



Automation-driven use case examples:

- Automated reconciliation of intercompany balances
- AI-based matching of bank entries to ledger transactions
- Extraction of filing deadlines from statutory calendars
- Contract clause extraction and risk flagging.

Each use case is scored for feasibility, value, and reusability. The goal is not to deploy everything, but to focus where AI can demonstrably enhance productivity and reduce operational friction.



GenAI workflow examples:

- Drafting of transfer pricing documentation sections (local and master files)
- Generating International Financial Reporting Standards (IFRS) disclosure narratives from financial data
- Drafting responses to tax authority queries with jurisdictional logic
- Summarizing new Value Added Tax (VAT) legislation by market
- Drafting budget and performance commentary for monthly packs

A platform that understands workflow, context, and change

The platform design rejects generic AI chat tools in favor of tailored workflows. Each use case is configured as an AI-powered process that guides the user through input, applies organisational knowledge, and produces actionable outputs.

What makes this platform resilient is its modularity? Each workflow can be:

- Independently maintained and improved
- Tied to version-controlled prompts and data sources
- Updated as regulatory or internal requirements evolve

New models can be introduced without redesigning the entire platform. Data sources can be extended as more content is structured and approved.



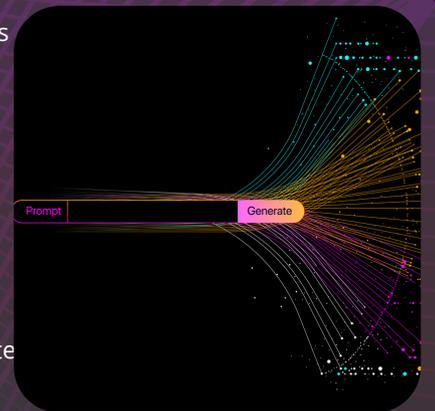
Data that is secure, governed, and strategically curated

No genAI platform is useful without content. However, dumping everything into a model is not strategic.

The organization invests in structured ingestion, tagging, and governance of key internal documents:

- Internal technical guidance and playbooks
- Jurisdiction-specific tax and regulatory content
- Past advisory memos and review checklists
- Audit schedules and enterprise resource planning (ERP) extracts

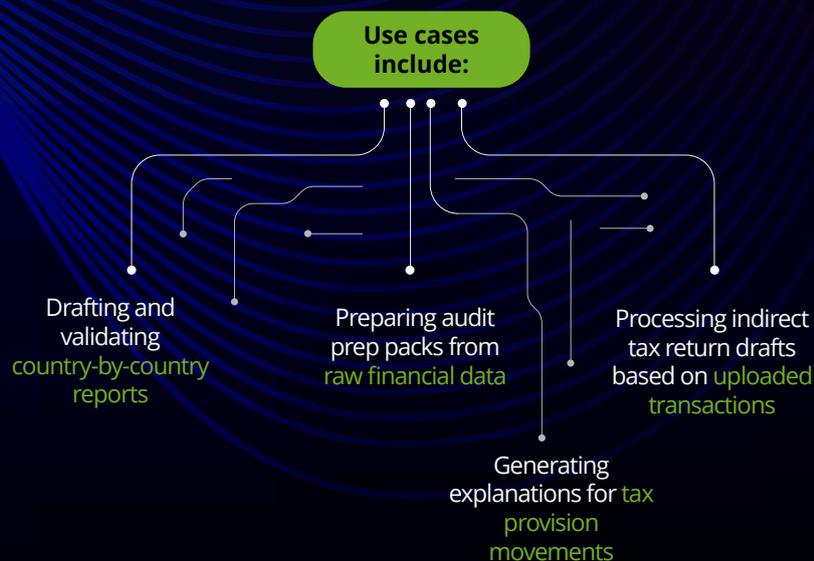
Content is organized by sensitivity. Redaction is applied where needed. Retrieval-augmented generation (RAG) infrastructure is designed to help the model pull only from approved sources.



A pilot that proves business value—not just technology

The pilot rollout targets two core departments: international tax and financial operations. It launches with thirty users and twenty workflows.

Success is measured not just by time saved, but by reviewer effort reduced, consistency improved, and knowledge reuse achieved.

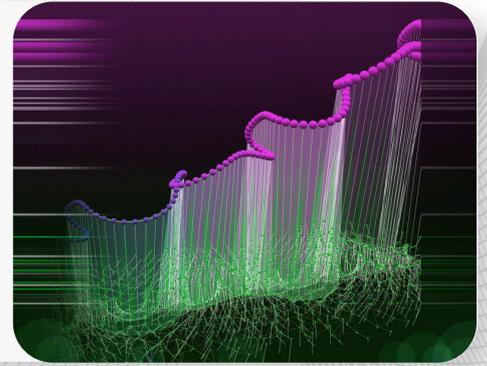


Governance that is designed in, not bolted on

From the outset, an AI governance framework is embedded:

- A genAI Steering Committee oversees workflow development and model testing
- Prompt libraries are curated and version-controlled
- Review protocols aim to have high-risk outputs receive human validation
- Usage data is monitored for compliance and drift

This supports the system's ability to grow safely—even as the technology landscape evolves.



Adoption that is human-centered

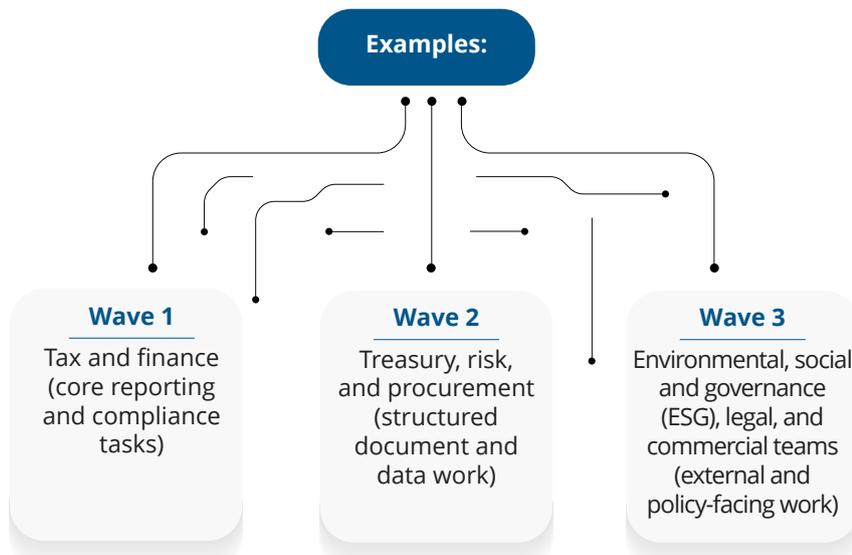


Training is tailored to roles. Learning content reflects specific workflows and tasks. Superusers and genAI Champions are appointed within each function.

Feedback loops are built into every interaction. Success stories are circulated weekly. Workflows are refined continuously. Performance data is visible to users and leaders alike.

Scalability that reflects demand, not ambition

The system does not scale all at once. Instead, a phased rollout plan is executed.



Each wave introduces ten–fifteen new workflows, tested and adapted before broader release.

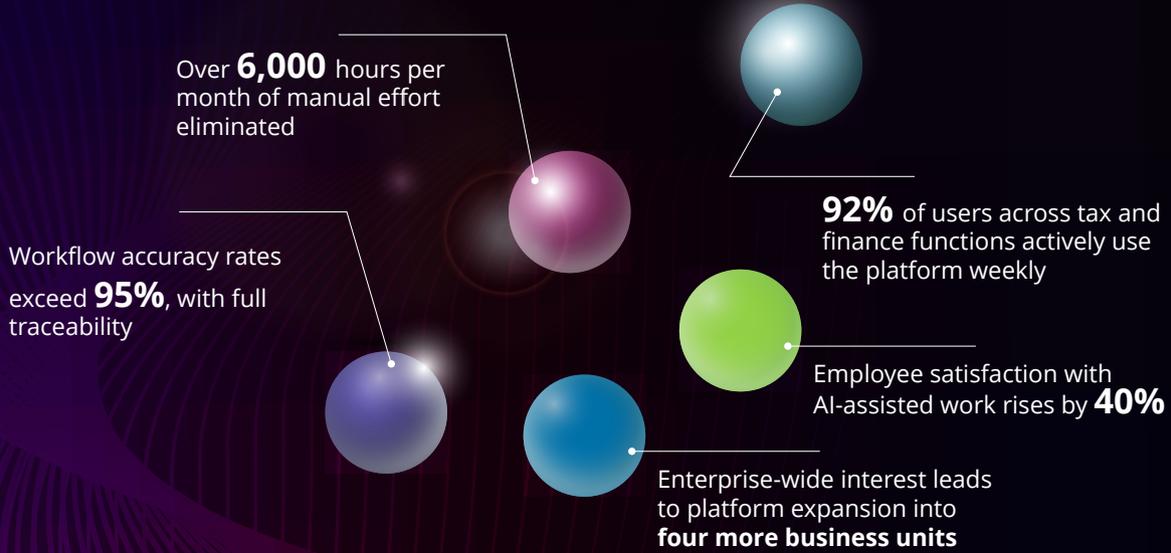
A strategy that is built for the future, not just the present

As new models emerge—faster, smarter, cheaper—the platform flexes. An architecture is established that supports model switching, new integrations, and adaptive retraining. AI is not hard-coded into legacy systems; it is layered modularly so that upgrades can happen with agility.

New governance challenges are anticipated. The genAI Steering Committee remains active. A quarterly model review cycle is instituted. Workflow libraries are refreshed. KPIs evolve—from usage to outcome quality to knowledge diffusion.

Outcomes that speak for themselves

After twelve months:



The performance outcomes cited in this case study are illustrative only and do not reflect results from any single Deloitte engagement. They are based on modeled projections and represent a synthesized vision of what could be achieved under optimal conditions using GenAI within tax and finance functions.

The final word

This future scenario story of successful genAI deployment is not just about technology—it is about design, governance, enablement, and above all, adaptability.

This organization will succeed not just because they moved fast, but because they moved deliberately. **They created a flexible AI architecture, grounded in real business use, structured around workflows, supported by humans, and built to evolve.**

In a space where hype fades fast, this is what strategic AI leadership looks like.

And it is already within reach.

Contacts



Mohamed Serokh
mserokh@deloitte.com



Joelle El Khoury
jelkhoury@deloitte.com



Ilham Hmimou
ihmimou@deloitte.com

Deloitte.

This publication has been written in general terms and therefore cannot be relied on to cover specific situations; application of the principles set out will depend upon the particular circumstances involved and we recommend that you obtain professional advice before acting or refraining from acting on any of the contents of this publication.

Deloitte & Touche (M.E.) (DME) is an affiliated sublicensed partnership of Deloitte NSE LLP with no legal ownership to DTTL. Deloitte North South Europe LLP (NSE) is a licensed member firm of Deloitte Touche Tohmatsu Limited.

Deloitte refers to one or more of DTTL, its global network of member firms, and their related entities. DTTL (also referred to as "Deloitte Global") and each of its member firms are legally separate and independent entities. DTTL, NSE and DME do not provide services to clients. Please see www.deloitte.com/about to learn more.

Deloitte is a leading global provider of Audit & Assurance, Tax & Legal and Consulting and related services. Our network of member firms in more than 150 countries and territories, serves four out of five Fortune Global 500® companies. Learn how Deloitte's approximately 457,000 people make an impact that matters at www.deloitte.com.

DME would be pleased to advise readers on how to apply the principles set out in this publication to their specific circumstances. DME accepts no duty of care or liability for any loss occasioned to any person acting or refraining from action as a result of any material in this publication.

DME is a leading professional services organization established in the Middle East region with uninterrupted presence since 1926. DME's presence in the Middle East region is established through its affiliated independent legal entities, which are licensed to operate and to provide services under the applicable laws and regulations of the relevant country. DME's affiliates and related entities cannot oblige each other and/or DME, and when providing services, each affiliate and related entity engages directly and independently with its own clients and shall only be liable for its own acts or omissions and not those of any other affiliate.

DME provides services through 23 offices across 15 countries with more than 7,000 partners, directors and staff. It has also received numerous awards in the last few years such as the 2023 & 2022 Great Place to Work® in the UAE, the 2023 Great Place to Work® in the KSA, and the Middle East Tax Firm of the year.