Tax for Growth

IGC

POLICY TOOLKIT

Harnessing AI and data for tax administration

Saugata Mittra, Twivwe Siwale, Laura Fras, Anders Jensen

Artificial intelligence (AI) and machine learning are rapidly reshaping economies and public institutions, including tax systems. For tax administrations, the key challenge is not whether to adopt AI, but how to create the right foundations to use it effectively.

This toolkit outlines why tax authorities need to establish governance frameworks for AI now, what institutional and data foundations are essential to maximise its benefits, and how AI can strengthen both enforcement and taxpayer services.

Mike Kononov @ Unsplash

JANUARY 2026

DIRECTED BY

LSE

UNIVERSITY OF OXFORD

FUNDED BY

UKaid from the British people

Contents

Why AI matters for tax administration

Artificial intelligence (AI) is booming. Industries worldwide are exploring the potential of AI, machine learning (ML), and big data.[1] Tools such as ChatGPT, launched only in November 2022, are seeing significant growth. As of August 2025, over 800 million weekly users utilise the platform (Duarte, 2025). AI has been recognised for its potential to reshape economies, restructure production, and generate value in ways previously unknown. The McKinsey Global Institute estimates that AI could generate an annual value of $2.6-$4.4 trillion across the global economy for the 63 use cases they examined (McKinsey, 2023).

While tax administrators are increasingly eager to leverage AI, the specific nature of revenue authorities presents distinct opportunities and challenges. Unlike private sector entities, tax authorities manage a unique ecosystem characterised by extremely rich administrative data, high-scrutiny decision environments, and massive operational scale. Implementing AI within such a complex framework – where decisions regarding audits and fraud impact millions – requires more than software. It demands that the foundational elements of data integrity, staff capability, and strategic vision be firmly in place before deployment. In an IGC survey of tax authorities in the Commonwealth Association of Tax Administrators (CATA), comprising 41 countries, 76% of respondents stated that they believe AI and big data analytics will bring about a full transformation to their tax authority. Beyond aims, 69% of revenue authorities said that, compared to five years ago, the usage of digital tools – precursors to AI – has significantly increased.[2] According to the OECD's Inventory of Tax Technology Initiatives, 48.5% of tax administrations actively use AI.

AI is already being used in various tax-related applications, from risk profiling to tackle evasion, to the use of virtual assistants in tax administration. HMRC's digital assistant had over 5.48 million interactions in 2024, and 62 per cent of those interactions did not require escalation to a human advisor (HMRC, 2025). HMRC estimates that the chatbot saved £5 million during the COVID pandemic (Trendall, 2023). A primary interest in AI is how it can build or significantly enhance tax capacity, particularly in developing countries where limited tax capacity has been a key constraint on tax revenue.[3] **Tax capacity is a tax authority's ability to collect taxes, at reasonable administrative cost, with limited economic distortions and low distributional impacts (Jensen et al., 2024).**

1 We refer to AI in this toolkit as a catch-all term for machine learning, big data tools and wider data analysis. For more information, we provide definitions below.

2 International Growth Centre and Commonwealth Association of Tax Administrations, *CATA 45th Annual Technical Conference - Pre-Event Questionnaire: Use of AI and Technology in Tax Administrations* (unpublished internal survey, July 2025). Full details of the methodology are presented in the Appendix.

3 IGC's work on "third-best policies" notes the severe informational barriers and significant enforcement constraints that developing country tax authorities face (Kleven et al., 2016).

This brief examines the current state of tax administrations' adoption of AI tools and offers insights on how they can position themselves to capitalise on the benefits of AI adoption. It draws on the results from an IGC survey of 25 Commonwealth countries polled on their use of AI, and the wider literature on AI. The evidence on AI remains nascent, and this brief presents a snapshot of the current state of affairs rather than an exhaustive evaluation of AI and its potential in tax administration. The brief is structured with four key messages.

- Building a strong foundation in data, talent, and governance is key to unlocking the potential of digital tools and AI.
- Digital technology, the precursor to AI, has been transformative for building tax capacity. AI similarly shows enormous potential to enhance tax enforcement and taxpayer services.
- Tax authority capabilities depend on institutional readiness.
- Sustainable integration of tools requires systematic implementation.

Box 1: What do we mean by AI?

Terms such as AI can easily become a catch-all term for any data science and analytics. There is substantial variation between different data processes. To distinguish between different technologies, we refer to the World Bank's Revenue Administration Handbook (Junquera-Varela & Lucas-Mas, 2024).

Artificial Intelligence (AI): is at its core concerned with building smart machines capable of performing tasks that typically require human intelligence. The term can also refer to a wide-ranging branch of computer science. An example of AI can include tax chatbots that use Large Language Models (LLMs).

Large Language Models (LLMs): are highly complex AI systems that learn from vast amounts of text data. Examples of LLMs include tools such as ChatGPT, which predict the most likely sequence of text (Chitty-Venkata et al., 2023)

Machine Learning: provides systems the ability to automatically learn and improve from experience by being trained on historical data. Examples can include methods such as Random Forest to detect characteristics of tax-fraudulent firms (Mittal et al., 2018).

Algorithms: a set of rules that precisely define a sequence of operations and are at the heart of computer science (Stone, 1971). Algorithms can have multiple different meanings depending on the situation. An algorithm in a tax context could cover anything from a simple line of code on a programming software to respond to emails, to more complex pre-filling of tax returns.

Big Data: a field that finds ways to analyse, systematically extract information from, or otherwise deal with data sets that are too large or complex to be dealt with by traditional data-processing software, such as Excel.

Data analytics: the science of analysing raw data to make conclusions about that information. In a tax setting, this could include a data analytics team that provides descriptive statistics of taxpayers.

Data visualisation: the graphical representation of information and data. By using visual elements like charts, graphs, and maps, data visualisation tools provide an accessible way to see and understand trends, outliers, and patterns in data. In a tax setting, an example of this would be a data dashboard.

> These terms can often become merged and used synonymously in colloquial settings. While there can be many overlaps in what these tools do, the data foundations and how they interact with each other can vary substantially by tool and tax authority.

This policy toolkit will primarily focus on the use of AI (including LLM chatbots) and ML for tax administrators. We will also explore tools such as data analytics and data visualisation, which allow a tax administration to make better use of the data it has and can be implemented at a much earlier stage of digital maturity. Increasing revenue depends on the enforcement capability of the tax authority, as well as voluntary compliance from taxpayers. Although both are complementary, the approaches required are vastly different. To make progress on both these levers, the tax authority needs a vision that encompasses data quality and completeness (as the fuel of the system), the skills capabilities of staff (as the engine of the system) and a culture of experimentation and constant improvement (a virtuous feedback loop). Having a framework to consider how to establish the foundations for integrating these frontier technologies into tax administration operations is crucial for making progress on digital transformation, improving compliance, and retaining citizen trust.

Success requires data, skills, and governance foundations

AI offers promising applications for tax administrations, but realising these benefits is far from automatic and depends critically on having the right foundations in place. Success requires a solid foundation in three areas: high-quality data, skilled staff, and robust governance frameworks for data privacy and security (Junquera-Varela & Lucas-Mas, 2024). These foundations become even more important as tax authorities adopt AI technologies.

Foundation 1: High-quality data

High AI capacity requires high-quality data. Digital tools and AI in particular require high-quality training data to produce useful outputs. AI, ML, LLM and big data models operate by training on large quantities of data to find statistical and probabilistic patterns, which are used to try to predict what will occur next. The more accurate the data, the more likely that AI will be able to model real patterns and predict outcomes that match reality as closely as possible.

Developing countries often struggle with high-quality data because of gaps in registration systems, large informal sectors, fragmented administrative records, limited budgets and technical capacity, and political or logistical barriers (World Bank Group, 2021). Tax, customs, social security, and business registries are often siloed, with no common identifier. That prevents merging data, increases duplication, and raises reconciliation costs for analysts (Bachas & Jensen, 2023).

AI and ML have the potential to make previously unusable or unreadable data accessible. By using tools such as natural language processing (NLP) and optical character recognition (OCR), tax authorities can digitise their archives of manual, paper-based data. AI tools can also help to improve the quality of data, for example, by automating data ingestion, standardisation and normalisation (by automatically mapping varied inputs to a standard format) or ensuring data consistency through fuzzy matching and probabilistic entity resolutions, which handles typos or variations, unlike traditional rule-based systems (Walz, 2025). The advantage of using these tools is their speed, efficiency, and much higher accuracy. This offers developing countries a relatively cost-effective way of mobilising their extensive administrative data to strengthen tax enforcement.

Box 2: What do we mean by data quality?

Data quality broadly refers to **whether the data is fit for purpose** (UK Government, 2020). The quality of data can be defined through six dimensions. For AI to be fully integrated into a tax system, high-quality data is needed to be fit for purpose.

Completeness: refers to the extent to which records are present. For example, if a customs tax data set has records dating back 100 days but is missing 2 days' worth of data, the data set will only be 98% complete.

Uniqueness: refers to the extent to which there is no duplication in the records. If you had a customs data set that was fully complete, but all the firms had the same taxpayer ID, the data set would have low uniqueness, so it would not accurately reflect the real world.

Consistency: the data does not contradict itself. For example, if a customs tax data set held 10 years of import data for a firm that only registered last year, this would not be consistent.

Timeliness: refers to the extent that data is an accurate reflection of the time period and that it is up to date. In a customs tax setting, untimely data might not include data for a firm for the last 2 years.

Validity: describes whether the data is in range and formatted as expected. A valid data set for customs tax would be one where the date of payment is saved in a date format and not a text format, and where the taxpayer ID is registered for customs tax and not only turnover tax.

Accuracy: to what extent the data matches reality. This is the hardest dimension to check and is the aim of a high-quality dataset.

The more data a tax administration can draw from diverse sources, the better its AI systems can detect patterns, assess risk, and generalise across different types of taxpayers. Greater volume and variety of data strengthen accuracy, robustness, and the potential for advanced applications like fraud detection and compliance prediction. Integrating across data sources opens the possibility for more complete and accurate data since different datasets capture different aspects of taxpayer behaviour. In the IGC-CATA survey, **36.3% of respondents stated that integration and interoperability of data was the top priority** when adopting digital tools.

Once datasets are integrated across departments, tools such as dashboards can be developed even before developing and deploying AI tools. For example, an integrated data set with a dashboard would allow an auditor to compare a firm's corporate income taxes with the salaries paid out to their employees, and the amount of VAT collected from e-invoicing or EFDs. In the case of Tanzania, a VAT discrepancies flag and comprehensive reports of taxpayers were created due to the integration of data from EFDs with taxpayer data (Siwale et al., 2025). This data integration and dashboard allows tax officers to cross-check customer receipts with seller invoices to spot anomalies and discrepancies, making auditing much more efficient. The key here is visibility. Without transactions being made visible through EFDs, revenue authorities could not hope to tax these transactions.

Most countries that implemented EFDs reported that firms found ways to circumvent the system, especially in an environment of weak enforcement. Furthermore, even when EFDs are in place, tax authorities in developing countries sometimes struggle to utilise the data they collect. Almunia et al. (2024) analyse Uganda's VAT returns data and find that sellers and buyers reported different transaction values in 79 per cent of firm-pair-month observations. This was not detected by the tax authority. The findings from such analyses should be paired with enforcement interventions such as inspections and audits.

Foundation 2: Analytical and modelling skills

Data analytics and modelling skills are the second crucial foundation for AI integration in tax systems. Without data analysis and engineering skills, the tax authority cannot respond to changes in tax codes, technology and taxpayers and is reliant on external support, an added cost that tax administrators in lower-income countries may not have budgeted for.

Dedicated data units within tax administrations have the potential to drive big changes in AI and technology adoption. In our survey, 27% of respondents already had a specialised AI or big data team and another 33% have plans to develop such teams. Such data labs operate in tax authorities in higher-income countries, such as the UK's HMRC or the Research, Applied Analytics, and Statistics (RAAS) division at the US Internal Revenue Service (IRS). We explore what this looks like below.

Create specialised data analytics teams

A range of data analysis and engineering skills is needed to digitise, integrate and manage the types of data used by most revenue administrations, a non-trivial requirement. Most units include some combination of the following:

Job title	Job description	Relevance to a tax administration
Data scientists and analysts	This role entails the ability to use data analysis and statistics to solve data-relevant questions. Data scientists often study the subject at university level.	In a tax administration, data scientists can help integrate, manage and clean data sets, run the data analysis that feeds data dashboards or run predictive models for different tax offices.
AI product manager	An AI product manager will work with data scientists to develop AI models and systems to apply AI tools to a relevant problem. They traditionally also come from a data science or AI background.	An AI product manager will lead the development of specific AI tools for tax purposes, such as an LLM for tax chatbots, ensuring they are fit for purpose.
Data engineer	A data engineer will design, build and maintain data systems. Data engineers will traditionally know how to code in different 'languages' (e.g. Java, Python, SQL).	A data engineer may build the systems used by data scientists, such as helping to build the platform for data dashboards or developing a data security infrastructure.
Data quality and metadata specialist	Maintains data standards and metadata, runs automated validation, and coordinates remediation to keep core datasets reliable; typically requires SQL, and proficiency in Python or R for rule-based checks.	A data quality specialist improves the accuracy of taxpayer and transaction records, reduces manual reconciliation, and increases the reliability of analytics and compliance models.
IT and cloud administrator	IT and cloud administrators manage on-site and cloud-based platforms, but they may not necessarily be the ones to initially build it (which a data engineer might).	If a tax administrator outsourced the creation of their cloud data management system, an IT and cloud administrator would be on-site to manage the system and make sure that it is integrated into daily operations.
Ethics and compliance leads	This role is needed to maintain the integrity of data security and be forward-looking in how data might be misused within a tax authority.	An ethics and compliance lead works on developing a data governance and ethical framework within the revenue authority.

These units produce user-friendly datasets that integrate data from across the tax authority (datasets linked to third-party reporting data), and create algorithms and codes that encourage greater usage of data and AI tools within the revenue authority.

Box 3: Data lab in Senegal

In 2015, the Senegalese tax administration (DGID) collaborated with researchers to create a data analytics unit or tax data lab (Bachas & Jensen, 2023).

Aim of the lab: A dedicated unit of six data scientists worked to digitise data, improve audit selection, detect tax evasion by firms, and enforce property tax. The unit created a standardised practice to digitise manually filed information.

How has the data lab improved outcomes?

It produced a comprehensive guide to navigate the complex information environment and suggested improvements to the IT departments and local tax offices. It also revealed data anomalies for which data checking solutions are progressively being implemented. Furthermore, it increased the administration's capacity to analyse and cross-check tax declarations and consequently detect non-compliant behaviours at a large scale.

Develop data skills throughout the tax administration

Using and feeding into a high-quality data infrastructure requires work from all who are part of the system. Beyond creating dedicated data units, **the whole tax administration needs to have a strong basis for using digital data management tools.** Any output using poor-quality data will be poor.

Data entry should be as simple as possible for on-the-ground tax collectors, with supporting training on how to use digital tools and software. Smart invoicing for VAT should make this process easier, as invoices will not require human entry, but no tax system can fully avoid human input on data. For example, if audits are done manually, any information obtained should be uploaded accurately to the data management system. **Data must not remain siloed in a personal Excel sheet or on paper.**

Training is vital. Staff should be trained to understand and use the outputs from a data unit, such as software and data dashboards. In Finland, the introduction of a new digital software saw over 4,000 personnel in the revenue authority trained on its use. Because the new system was well integrated and used by most personnel, it led to 70 legacy systems being decommissioned and 15-20 million euros saved in IT costs (OECD 2020).

Data systems and AI are technically demanding, so tax authorities face a clear trade-off: building in-house creates ownership, keeps sensitive data inside the administration, and gives the authority the flexibility to adapt, correct and evolve systems through pilots that build technical skills; but it also requires sustained investment in staff and capability.

Outsourcing – or a hybrid model combining vendor expertise with in-house oversight – is common and pragmatic where talent is scarce: our survey found 73% of respondents used a hybrid approach. Vendors bring specialised AI knowledge and speed, but to realise long-term benefits, authorities must retain a capable data unit to provide governance, ensure solutions are fit for purpose, integrate tools into workflows, and capture skills through joint project development.

Foundation 3: A governance and ethical framework

Tax authorities hold some of a country's most sensitive data, so any AI or big-data rollout must be anchored in a clear governance and ethics framework aligned to the authority's domestic resource mobilisation goals; 53 per cent of respondents in our survey have already begun formal AI strategies tied to DRM, and many rank compliance and security among their top integration concerns. The framework should set out why AI is being used, which programmes or workflows it will support, and the expected benefits, while balancing those gains against risks from data misuse and erosion of taxpayer trust.

To build trust and ensure responsible use of AI, every tax administration should develop a clear governance and ethical framework. This framework rests on three core areas: **privacy and ownership, data security, and data use and transparency.**

Privacy and ownership mean being explicit about who can access data, for what purpose, and how personal information is protected. Tax administrations should be transparent about why data is collected, how long it is kept, and how anonymity is ensured. For example, the UK's Information Commissioner's Office and OECD guidelines emphasise the importance of defining clear responsibilities for data access and use *(Guidance on AI and Data Protection, 2025; OECD Legal Instruments, n.d.).*

Data security is about protecting the systems and information that underpin AI tools. Administrations need clear rules on how data is stored, who can view or share it, and how risks such as breaches or unauthorised use are handled. Simple but consistent safeguards - such as access controls, regular reviews, and clear accountability - are essential. The Australian Tax Office, for instance, publishes assurances that AI tools are reviewed regularly and that all taxpayer-affecting decisions remain under human supervision *(ATO AI transparency statement, n.d.).*

Data use and transparency relate to how AI models are applied in practice. Policymakers should make sure that every AI project has a clear purpose, measurable benefits, and a way for taxpayers to understand how AI influences decisions. Publicly explaining what data is used, what the system does, and how people can question or appeal outcomes helps sustain public confidence. The UK's Algorithmic Transparency Standard provides a straightforward template for such communication *(Algorithmic Transparency Recording Standard Template, 2025).*

Finally, **human oversight** is non-negotiable. AI can support, but should not replace, human judgment in areas that directly affect taxpayers' rights or liabilities. Assigning staff as data or model "stewards," requiring review points before actions are taken, and maintaining clear appeal routes all help ensure accountability. The EU's new AI Act and OECD's recommendations both stress this balance between innovation and human control *(The AI Act, n.d.; OECD Legal Instruments, n.d.)*.

When these principles are in place – clear responsibilities, secure data, transparent use, and routine human oversight – AI can enhance efficiency and fairness while protecting taxpayer trust and administrative integrity.

Saugata Mittra,
Policy economist, IGC

> Transitioning from theory to application requires a shift in perspective: AI adoption is a pathway of incremental gains. The focus should remain on selecting tools that fit your specific institutional context. The question is not if you should adopt AI, but which tools match your current institutional reality. Success lies in aligning your ambition with your readiness - ensuring your internal 'engine' of data and governance is strong enough to power the external applications you want to deploy.

Digital transformation drives tax capacity gains

Once these foundations of data, skills, and governance are established, administrations can look to the history of digital transformation to understand how to apply them. Ensuring the right foundations are in place allows tax administrations to tailor their tools to their needs. Over the past two decades, technology has been the primary tool that tax authorities have used to transform tax administration. Many countries have invested heavily in technology, digitalising tax functions and rolling out tools such as Electronic Fiscal Devices (EFDs) and e-invoicing programmes.

In our survey of Commonwealth countries, more than half of the countries state that over 50 per cent of their processes are now automated. This compares with about 38% of respondents to the OECD's Tax Technology Initiatives survey on whether robotic process automation is actively used in the administration (OECD, n.d.). Forty-four per cent of IGC's respondents indicate that, compared to five years ago, the use of digital tools in their department has changed significantly. Since 2018, the number of tax administrations integrating AI has almost doubled, going from 26% of administrations to 48% in 2023, according to the OECD.

The rollout of digital technology in tax administration can reveal a great deal about the potential of AI to reshape tax administration. Historically, the biggest leaps in improving tax capacity have come from the use of **third-party reporting** and interventions that have targeted **tax evasion**.

Third-party reporting

Third-party reporting occurs when information about economic agents is made directly available to the tax authority by third parties. It was the key driver that led to a significant increase in tax revenue from 15 per cent of GDP to 40 per cent of GDP in many of today's developed countries (Jensen, 2022). **The essential role of third-party information is to make taxable economic activity visible to the tax authorities.** What can't be seen can't be taxed.

One of the best examples of third-party reporting is when employers report the income of employees directly, for example, through payroll withholding. A study in Denmark, which provides evidence for the effectiveness of third-party reporting, found that tax evasion on employer-reported income was close to zero. However, for self-employed workers, the rate was 44.9% demonstrating the effectiveness of third-party reporting (Kleven et al., 2011). The value-added tax (VAT) is another example of the importance and value of third-party reporting. With the VAT, buyers and sellers have different incentives in reporting taxable activity. Buyers want to claim the VAT due on purchases, and sellers would prefer to underreport their sales. When a buyer submits a VAT return to the tax authority to claim their VAT, they provide third-party information about the seller, which ensures that the tax is remitted.

Technology and digitisation have led to a substantial increase in third-party information and its accessibility to tax authorities. Waseem (2025) argues that rather than increasing third-party information, technology has made it more accessible to tax authorities, but questions whether it is being utilised. **One potential advantage of AI is its capacity to analyse large volumes of data that tax authorities now collect by virtue of digitalising their systems.** Tax administrations are increasingly in the business of information management. Given the diversity and variety of administrative and tax data collected, using digital tools and AI has increasingly become a necessity for tax administrations.

Targeting evasion

The second major factor that leads to increases in tax capacity is effectively increasing the probability – and the perceived probability – that tax evasion by economic agents will be detected by the tax authority. The effect of third-party reporting is heightened when firms know that the tax authority has information on their activities and follows through with enforcement. Reducing evasion is one of the best ways of adding to the revenue pool.

Using a well-established model in the public economics literature of tax evasion, we can decompose the probability of a conviction for tax evasion as the probability that a case is selected for audit multiplied by the probability that if selected, evasion is actually detected by the tax authority, multiplied by the probability that the case is convicted given evasion was detected (Allingham & Sandmo, 1972). AI increases the probability of evaders being selected for audit with better selection of cases (based on analysis of patterns and trends in data) and risk-scoring and classification of taxpayers or goods. AI also significantly increases the probability of detecting evaders, given that they have been selected for audit due to much faster and better evidence extraction and data matching.

The capacity of AI to produce more third-party information, or to improve the ability of tax administrators to act on this information, will determine its potential for tackling tax evasion and improving compliance. AI as an enabler to catalyse efficiency in traditional methods of tax compliance is a powerful frame of reference. AI not only helps generate richer, more complete third-party information through advanced analytics and integration, but also enables tax administrations to interpret, prioritise, and act on this information far more efficiently.

What we know about digital tools before AI

E-filling, e-registration, electronic fiscal devices (EFDs) and e-invoicing are some of the most ubiquitous technologies that tax authorities have adopted. The process of evaluating the effectiveness of digital tools varies widely. Most tax authorities would prefer cost-benefit analysis or calculations on the rate of return on investment. An economic evaluation is another useful approach. This involves conducting empirical studies to determine the implications of a technology rollout on economic measures, such as revenue growth, compliance, and enforcement. The economic evaluation literature is reviewed here.

EFDs have been widely adopted in many developing countries. EFDs send real-time transactional data on firms' sales activity directly to the tax authority. They also produce a standardised invoice that is at the heart of VAT compliance. EFDs can address problems of low compliance by increasing the (third-party) information available for tax authorities to enforce VAT. The best evidence for the effectiveness of EFDs is from Rwanda, where an IGC evaluation found a 5.4 per cent increase in VAT revenues (Eissa & Zeitlin, 2014). However, Kenya and Zambia did not see any increases in revenue following the rollout, and in Zambia, preliminary evidence shows a decline in revenue (Dillon et al., 2023). Contrastingly, Ethiopia saw increases of 12% for income taxes and 48% for VAT after introducing the equivalent of EFDs (Mascagni et al., 2021). The mixed evidence shows that EFDs are not a complete solution.

Current digital compliance measures can only automatically detect unilateral evasion – when one party reports a transaction and the other does not (or reports a different amount). The other challenge can be distinguishing between genuine mistakes and fraud. AI might be better placed to address both of these issues by utilising neural networks to graphically represent links between entities in the supply chain, detecting VAT carousels (Vatcalc.com, 2025). But the nature of VAT means that basic digital cross-checks can detect simple mismatches before integrating complex AI/ML models. Some empirical studies have shown large revenue dividends from (pre-AI) digital tools. For example, in China, e-invoicing transformed VAT collection and contributed to a 27.1% increase in VAT revenues (Fan et al. 2018).

Bellon et al. (2022) use administrative data on over 53,000 firms in Peru to show that the introduction of e-invoicing increased reported firm sales, purchases and value-added by more than 5% in the first year, with effects growing to 10-15% by the fourth year after adoption.

Box 4: Catching "Bogus Firms" in India

What happened: In the National Capital Territory of Delhi, tax authorities faced a problem with "bogus" (shell) firms. These entities existed only on paper, registered solely to issue fake invoices that allowed other companies to claim fraudulent VAT credits. This scheme was estimated to cost the government roughly **$300 million annually** in lost revenue.

The AI solution: Researchers collaborated with the tax administration to deploy a ML model trained on years of VAT return data. Instead of looking at firms in isolation, the AI analysed the **network of transactions** (who sold to whom), identifying complex patterns of "circular trading" and impossible profit margins that human auditors could not easily spot (Barwahwala et al., 2024).

Results and Key Lessons: Technically, the AI was a triumph: it successfully identified bogus firms with high accuracy, uncovering an estimated $15–40 million in potential annual revenue recovery. However, the field evaluation revealed a critical insight for administrators: despite the model's accuracy, actual revenue collection did not significantly increase. The bottleneck wasn't the AI, but the operational workflow – auditor teams faced institutional friction and resource constraints that prevented them from effectively acting on the new intelligence.

The case study illustrates that AI can accurately identify potential revenue, but it cannot collect it. Success requires integrating AI tools directly into staff workflows and aligning incentive structures to ensure acting on AI insights is feasible for human officers.

AI enhances enforcement and taxpayer services

AI is the next stage of the digital transformation trajectory in tax administrations. This section provides five examples of the applications that AI can transform in revenue administration. We also provide a ranked list of AI applications by their likely return on investment for low-income countries.

Use case 1: Detecting evasion

Fraud and evasion can be countered with more and better information on taxpayers. AI can excel when used for detecting fraud and tax evasion, and typically performs better than individual tax officials due to the amount of data that AI can analyse. This can unearth patterns of fraudulent behaviour, enabling tax authorities to enforce tax compliance.

For example, customs officials are often overwhelmed by the sheer quantity of shipments to process. Case management and risk-rating channels may result in a significant number of shipments being allocated for manual inspection, where limited tax is at stake, thereby taking up precious time and attention that customs officials can scarcely afford to waste. AI excels in finding patterns in large amounts of data, making customs an ideal environment to deploy algorithms to limit the caseloads of customs officials and better flag high-risk shipments.

IGC-funded research in Paraguay is exploring how machine learning can significantly increase the number of customs violations detected from a smaller caseload, indicating much greater accuracy than the status quo (Finan et al., 2025). Another IGC study in Uganda examined how the installation of fixed cargo scanners at trade facilitation points significantly reduced the amount of misreporting at customs since data from these scanners is matched with corresponding trade documents (Okwany, 2025; Davies et al., 2024). The success of these examples stems from the tax authority identifying a specific problem they want to resolve (in Paraguay, the poor accuracy of the existing risk management system; in Uganda, perennial underreporting and misreporting at customs) and deploying a small, targeted solution.

AI has also been used to tackle VAT fraud – for example, in the case of ghost firms in Ecuador (Carrillo et al., 2023) and in Delhi (as described in Box 4). Ghost firms exist for the purpose of issuing fake invoices to firms, allowing these firms to claim input VAT for costs never incurred. Ghost firms are difficult to track, as they often have no permanent address and can easily close and reopen as a different operating entity. AI has been effective in identifying the *clients* of these ghost firms, making them an ideal candidate to target for enforcement activities. This has had knock-on implications for the VAT supply chain – because ghost firm clients were informed they may be audited, they started asking downstream suppliers in *their* supply chain to provide invoices. This enhances VAT's self-enforcing mechanism. The total amount of additional taxes filed increased by 81% compared to before the intervention.

AI can enable the detection of fraud and evasion at scale, whether applied in VAT or customs. Given how relatively new these technologies are, AI applications in this space will continue to evolve.

Use case 2: From descriptive to predictive analysis

A key strength of AI is its ability to shift processes from descriptive to prescriptive, going from "what happened" to "what will happen". As tax authorities boost their data and AI capabilities, they can shift from preparing descriptive, reactive analyses to predicting and preparing for future interventions which raise revenue. Descriptive data can include generating reports, statistics, or undertaking an analysis of what has already passed. For example, it could include the creation of a data visualisation of how much of each tax type has been paid in a given month, or how many firms have been audited within the last year, or EBM usage within a country. This type of data analysis is a helpful snapshot for revenue authorities and plays a foundational role in policy decisions within revenue authorities.

Predictive data analysis can include undertaking analysis that can forecast or attempt to predict how policies or economic climates may impact taxation, making tax authorities more proactive and less susceptible to playing catch-up. Examples of predictive data can include forecasts of how certain businesses may be impacted by a change in tax thresholds. Several revenue authorities (for example, the Australian Taxation Office, HM Revenue & Customs, and Receita Federal do Brasil) are combining e-reporting streams and machine-learning models to run scenario simulations and near-real-time forecasts that predict how policy changes or economic shocks will affect revenue and taxpayer behaviour, enabling more proactive, targeted policy responses (Australian Taxation Office, 2025; HMRC, 2024; *An Overview of Digital Transformation at Receita Federal Do Brasil*, n.d.). While predictive data analysis does not offer certainty, it is the next step in the data maturity process for revenue authorities. As revenue authorities expand their data capacity, predictive data analysis will become increasingly effective. Tax authorities move from using historical data to look backwards to using real-time data to better understand how taxpayers *respond* to the system.

Use case 3: Tax form processing & internal support

Tax offices must sift through high volumes of forms, letters, and documents daily. Manual review is slow, error-prone, and strains limited staff capacity. Large language models (LLMs) excel in this type of task, streamlining the routine processing of documents significantly. For example, LLMs can:

- Auto-fill forms & suggest deductions
- Summarise documents (e.g. contracts, receipts)
- Draft responses to taxpayers
- Interpret and summarise legal changes

Based on the World Bank's tax compliance cost surveys, small firms often face very high, regressive compliance burdens, with some estimates indicating up to 15% of turnover for the smallest firms (Coolidge, 2010). Recent IGC research (Waseem, 2025) bolsters these findings by showing that digitalisation can have disproportionately regressive effects for small firms, with these firms spending over 100% of VAT remitted on complying with new comprehensive transaction-level reporting requirements.

Incorporating AI tools can have a positive impact on revenue collection by reducing the cost for taxpayers to comply with their obligations or by reducing the cost of collection for the tax authority. AI can reduce the costs of compliance through applications such as automated document processing, which can automatically read invoices, receipts and bank statements and map them into tax categories, reducing or eliminating the need for manual entry. AI can be used to incorporate this data into prefilled tax returns, further reducing the time businesses spend complying with the law. This means fewer errors, less active recordkeeping for small businesses, and can also free up tax staff time for high-value activities that directly lead to revenue recovery instead of cross-checking basic data entry.

Use case 4: Knowledge management & training

Tax administrations hold vast institutional knowledge, from legal interpretations to procedural know-how, yet much of it is scattered across PDFs, outdated manuals, or held only in the minds of experienced staff. Chatbots can be an effective way of mobilising this knowledge for the benefit of the authority. AI chatbots use natural language processing to interpret a taxpayer's question, map it to the correct rule or knowledge-base response, and deliver a guided answer or escalate complex cases to a human agent. Just as chatbots can be the external face of knowledge management for tax authorities with taxpayers, LLMs can also act as internal tutors and librarians. TaxGPT, a private tool trained on US tax law, answers specific tax questions with legal citations. Tax authorities can similarly train an LLM on the corpus of corporate and income tax law, allowing tax office staff to ask questions like "how do I handle unreported VAT?" and get instant guidance, based on the law.

Using LLMs in this way would centralise access to rulings, manuals, circulars, and FAQs in natural-language format. It would provide on-demand learning during onboarding for new staff or when executing a task. Furthermore, it can capture and preserve expertise from retiring staff or siloed teams. The Rwanda Revenue Authority is exploring this application through a partnership with the African Institute for Mathematical Sciences (AIMS), potentially revolutionising the way new staff are inducted and taught on-the-job (RRA, 2024).

According to OECD analysis, staff costs dominate operating budgets, with salary alone comprising approximately 72% of tax authorities' operating spend (OECD, 2025). In addition, the cost of replacing an employee can range from approximately 16% to 200% of spending on annual salaries, depending on the education and experience required for the job (Abner & Hur, 2023). Poorly trained staff can also mean inconsistent information shared with taxpayers, creating real and lasting impacts on taxpayers' lives. Embedding LLMs would enable "always-on" coaching and access to policy memory, potentially improving training quality and reducing staff turnover as menial tasks are increasingly done by AI. LLMs can bridge knowledge gaps and build institutional resilience, ensuring every officer has access to the best guidance, 24/7.

Use case 5: Policy analysis & scenario modelling

As tax policy units within tax administrations know, policy complexity is a key issue in tax. Designing tax reforms requires a deep understanding of how laws interact, who is affected, and what precedent exists. However, legislative drafting is labour-intensive, and analysis is often reactive.

LLMs can enhance strategic policymaking significantly by:

- Allowing comparison of different reform options (e.g., changing thresholds or incentives) using international examples and benchmarks.
- Draft amendments using consistent legal language and flag conflicts with existing statutes.
- Summarise stakeholder input from consultations or public forums.
- Translate proposals into plain-language briefs for non-expert audiences.

LLMs can dramatically speed up fiscal policymaking when integrated with agent or microsimulation models to evaluate and explain fiscal policy choices (Karten et al., 2025). Microsimulation models are a tool that applies tax, benefit, and policy rules to individual- or firm-level microdata to estimate the distributional and fiscal effects of policy changes.

LLMs can translate policy text into model-ready rules for microsimulation engines (e.g., "raise the small-business VAT threshold to X"), reducing the time needed to implement scenarios and enabling analysts to run more policy variants quickly. Furthermore, LLMs can autogenerate multiple parameter sweeps (threshold levels, phased implementations, combinations of incentives) to produce a summary of each scenario and feed results into a microsimulator for numeric comparison. By synthesising international experience and suggesting realistic parameter ranges for each country, LLMs can help analysts test those benchmarks against local data. After a microsimulation run, an LLM can produce readable, stakeholder-tailored narratives explaining who wins and loses, whether a reform is progressive or regressive, and highlight trade-offs between revenue and equity. Finally, LLMs can speed up sensitivity and robustness checks by orchestrating scenario-based runs (with varying elasticities, compliance rates, and behavioural responses), flagging where policy conclusions depend on questionable or fragile assumptions.

In Canada, the revenue service uses LLMs to route customer feedback to the appropriate department using sentiment analysis – understanding the tone of a message and devising appropriate responses for the case officer. In Bangladesh, LLMs have been used to help ensure new SROs align with existing law and avoid duplication or conflict (Wasi et al., 2024). Among currently available open-access tools, Claude (Anthropic) is best suited for policy work due to its ability to process long legal texts, draft amendments, and translate complex tax reform into plain language.

LLMs support smarter, faster, more inclusive policymaking by serving as analysts, editors, and translators.

Highest Return on Investment Use Cases for Low-Income Administrations

Recognising the budget and capability constraints for administrations in low-income countries, the following use cases of AI provide the greatest return on investment and should be prioritised when first building capabilities in this space:

Rank	Area	Why High Payoff
1	VAT/GST Fraud Detection	Largest leakage; network fraud detectable via AI (Alexopoulos et al., 2025; Rhode, 2025; Erastov & Balytska, 2025)
2	Customs & Excise Fraud Detection	High leakage in import-dependent LICs (Junquera et al., 2025; Kelleher, 2025)
3	Property Tax & Geospatial AI	Large off-roll base; satellite/ML expands coverage (Junquera et al., 2025; Blanco, 2024)
4	Domestic Audit Risk-Scoring	Improves audit yield where capacity is scarce (Eberhartinger et al., 2021)
5	Behavioural Compliance Models	Low-cost, improves voluntary compliance (Blanco, 2024)
6	Transfer Pricing / Trade Mis-Invoicing	Relevant where some MNE trade exists (OECD, 2025)
7	High-Wealth Individual Detection	High yield but a limited number of cases (Pickett, 2025)

As explored later, gains from evasion and application of new technologies should be reinvested in developing more advanced capabilities and securing political support for long-term reform of the tax authority.

Capabilities depend on institutional readiness

Having a framework to think about what tax authorities can achieve with their current level of digital capability is crucial. It allows countries to benchmark themselves against peers, understand best practices, and concentrate efforts on dimensions of the tax authority's operations that enable more advanced technology to be effectively deployed.

Ultimately, what the tax authority can achieve is determined by its level of digital maturity. The World Bank defines a tax administration's level of digital maturity with respect to three dimensions: scope, data quality, and type of use (Junquera-Varela & Lucas-Mas, 2024). The following definitions come from the Handbook for analytical consistency.

Scope refers to the level of integration of data analytics in an organisation – at the most basic level, very few individuals, if any, are using advanced data analytics in isolated parts of the administration. At an intermediate level, localised groups use advanced data analytics in specific areas of tax administration. At a more advanced level, the use of analytics is generalised across the administration by all departments and for all core business processes and functions.

Data quality refers to the quality of data available in the tax administration's repositories. Data analytics will provide poor results if trained on low-quality or incomplete data. Refer to Box 2 for details on how data quality is measured and defined.

Type of use refers to whether the administration uses advanced data analytics to find out more about what happened in the past (descriptive) or whether it uses this data to reshape its operations for the future (predictive).

Digital maturity is like a ladder – each step is sequential and builds the necessary capabilities for the next step, but data quality is the fundamental binding constraint. The quality and scope of available data largely define the level of digital maturity in a tax administration.

Figure 1: The Digital Maturity Framework

Measuring progress across the three dimensions of institutional readiness.

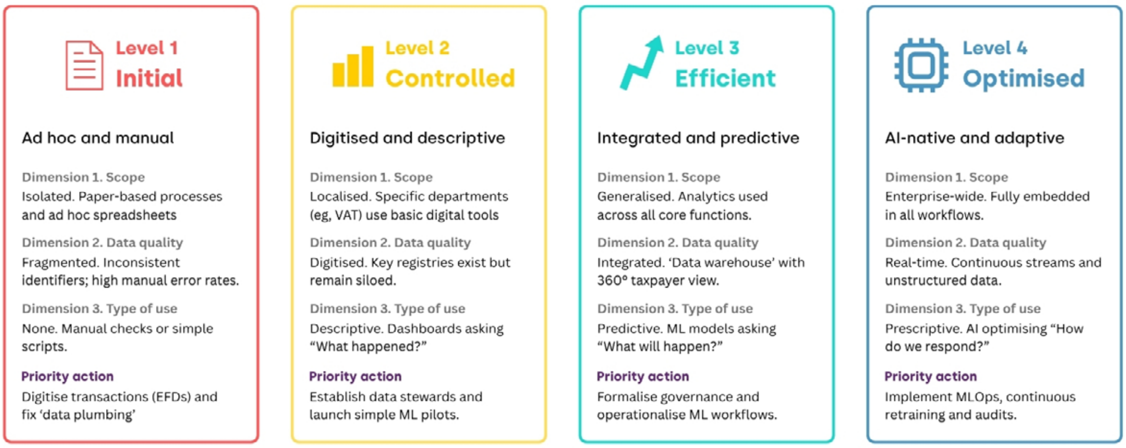

Improving the quality and scope of data is therefore the initial priority. Digitising transactions (EBMs, e-filing) and basic data integration are important first steps. Collaborating with external partners (e.g., CATA, universities, the World Bank, the IMF, etc.) can build internal data analytics capacity by leveraging tools such as ML models. For more advanced tax authorities, investing in e-invoicing and digital reporting further increases the scope and quality of data.

Digital maturity advances through iterative cycles: a tax administration builds a digital tool or predictive model, the prototype exposes gaps in data and processes, and those gaps are then cleaned and integrated back into the system. As data quality improves, model performance improves, creating a virtuous loop that makes tools more reliable and useful. Crucially, this cycle also develops in-house capability. Repeated pilots and hands-on refinements let staff acquire practical AI and machine-learning skills through learning by doing. See the Appendix for a checklist to assess digital maturity within your tax administration.

A good way to get started is to find specific, simple use cases. These starting points can measurably improve revenue recovery while linking to a broader strategic plan on utilising AI effectively in your local context. Many of the most successful initiatives take small incremental steps. Simple systems with modest goals can have significant payoffs over time. Please refer to the previous section for a table on the highest ROI use cases for AI in tax administrations in low-income countries.

Conducting pilots and small experiments is crucial for building an understanding of where these tools might be most effective, the level of capability in the revenue administration, and political support. Investments in digital maturity will sometimes see returns accrue over decades, showing the importance of keeping political leadership onboard to sustain momentum. Improving the data infrastructure of the tax authority by setting it up for the future involves building flexibility directly into the way the system operates. Having a framework helps to identify the non-negotiables – the specific steps that cannot be compromised to advance digital maturity. The World Bank's GovTech Maturity Index 2022 notes that income level is highly correlated with digital maturity. 65% of GovTech leaders are from high-income countries. 22 out of 28 low-income groups have minimal focus on GovTech, but this is not causal (GovTech Maturity Index: 2022 Update - Trends in Public Sector Digital Transformation, 2022).

Increasing levels of digital maturity within a revenue administration should be seen as a key enabler for the core policy objectives of the tax authority. It enables more advanced and sophisticated methods to be used for specific objectives, as outlined in the table below. For boosting voluntary compliance, for example, increasing digital maturity allows tax authorities to shift from using virtual chatbots to answer taxpayer queries to being able to accurately predict tax returns for self-employed individuals. As the tax authority builds capabilities and momentum with a specific tool, new instruments for increasing domestic resource mobilisation become available.

Objective	L1 ("initial")	L2 ("controlled")	L3 ("efficient")	L4 ("optimised")
Boost voluntary tax compliance	Provide basic digital information such as downloadable forms, simple FAQ pages; introduce SMS or email reminders for filing and payments	Use virtual agents to reply to taxpayers' questions (Zambia)	Sentiment analysis of taxpayer/trader feedback (Canada)	Predict accurate tax returns for self-employed and sole proprietors (Italy)
Strengthening risk management and segmentation	Perform manual segmentation using spreadsheets (size, sector, geography) to support basic risk triage.	Assign probable economic activity to unclassified taxpayers (Australia)	Automatic risk profiling of taxpayers and new traders (Spain)	Discern multilayered relationships among traders/taxpayers (Italy)
Detecting and preventing fraud	Apply simple rules based on thresholds or inconsistencies, for example flagging missing TINs or mismatched amounts between declarations and payment slips	Singular outlier detection in declarations and refund requests (Serbia)	Risk-group classification for declarations or VAT filings (Brazil)	False document spotting, import/export fraud prediction, tax-crime propensity scoring of individual taxpayers (India, Mexico, Colombia)
Optimise operational efficiency and resources	Implement basic workflow tracking such as simple case logs for registration, audits or refunds	Resume screening and bias detection for hiring (UK)	Process-drift detection in core workflows (Rwanda)	Automatically process suspense accounts (Australia)
Improve audit and collection effectiveness	Compile consolidated taxpayer case files combining registration, past audits and payment history	Estimate cost and probability of arrears recovery (South Africa)	Image analysis of containers/baggage (Uganda)	Automated valuation of goods in customs (China)

Policy Recommendations

This toolkit has argued that AI adoption in tax administration is not a technological compliance formality, but a fundamental shift in how tax authorities manage information. To navigate this transformation, we propose a dual-track framework: building the **Internal Engine** (capacity, data, and governance) while simultaneously deploying the **External Application** (enforcement, service, and policy).

As illustrated in the figure below, these two tracks are not separate – they form a virtuous feedback loop. The engine powers the application, and the revenue gains from the application provide the political capital and resources to upgrade the engine. While the recommendations focus on lower levels of digital maturity, the same framework will apply, irrelevant of your starting point.

I. Building the Engine: Internal Capacity

The "Engine" represents the institutional foundations required to operate AI safely and effectively. Without this, external tools will fail or create unacceptable risks.

1. Improve the data infrastructure first: AI models are only as good as the data they are trained on. Before procuring advanced AI tools, administrations must ensure their data is "fit for purpose" (see Box 2).

   a. Prioritise the digitisation of transactions (e.g., through EFDs or e-invoicing) over the procurement of complex predictive models.

   b. Establish a "single reference point" by mapping all legacy identifiers to a unique taxpayer identification number (TIN).

   c. Integrate data siloes (e.g., merging customs, VAT, and income tax registries) to create the 360-degree taxpayer view required for effective risk-scoring.

2. Build the "human-in-the-loop": technology cannot replace the context-specific expertise of tax officials. The goal is to reduce friction between technical outputs and frontline decision-making.

   a. Establish a small, in-house data analytics unit that acts as a bridge between IT vendors and tax operations. This unit ensures ownership of the models remains within the administration.

   b. Mandate "human-in-the-loop" protocols where AI outputs (such as audit flags) are treated as recommendations requiring human verification, rather than autonomous decisions.

3. Establish the "governance framework": public trust is essential for tax administration. A lack of clear governance can lead to privacy breaches or algorithmic bias that erodes voluntary compliance.

    a. Implement a governance framework that includes Privacy Impact Assessments (PIAs) for every new AI pilot.

    b. Define clear accountability lines – designating specific "Data Stewards" responsible for the quality and security of specific registries.

II. The Application: External Impact

The "Application" is where the engine is applied to real-world problems to achieve domestic resource mobilisation (DRM) goals.

1. Sharpen enforcement: AI excels at identifying non-linear patterns that human auditors miss. Administrations should target high-leakage areas where AI offers the highest return on investment.

    a. Move from random audits to risk-based selection, using predictive scoring to allocate scarce audit resources to the highest-risk cases.

    b. Deploy machine learning pilots in VAT fraud detection (using network analysis to identify carousel fraud) and Customs (using risk-scoring to detect undervaluation).

2. Streamline compliance: digitalisation can be regressive for small firms if compliance costs are high. AI should be used to lower the barrier to entry for the formal sector.

    a. Utilise LLMs to power chatbots that can answer taxpayer queries 24/7, standardising advice and reducing call-centre loads.

    b. Use data integration to pre-fill tax returns, significantly reducing the time and cost of compliance for compliant taxpayers.

3. Optimise policy design: AI allows administrations to move from descriptive statistics (what happened?) to predictive modelling (what will happen?)

    a. Integrate LLMs with microsimulation models to "stress test" proposed tax reforms against different economic scenarios before implementation.

III. The Feedback Loop: Sustaining the Cycle

Sustainable transformation requires an iterative approach. Administrations should not attempt a simultaneous, large-scale implementation. Instead, they should:

1. Start small: launch targeted pilots (e.g., a single VAT risk model) to demonstrate value.
2. Reinvest gains**:** use the revenue recovered from early enforcement pilots to justify budget increases for data infrastructure and staff training.
3. Iterate: continuously retrain models as data quality improves, moving up the Digital Maturity ladder from simple descriptions to advanced predictions.

Figure 2: The "Engine & Application" Framework for Tax Administration AI

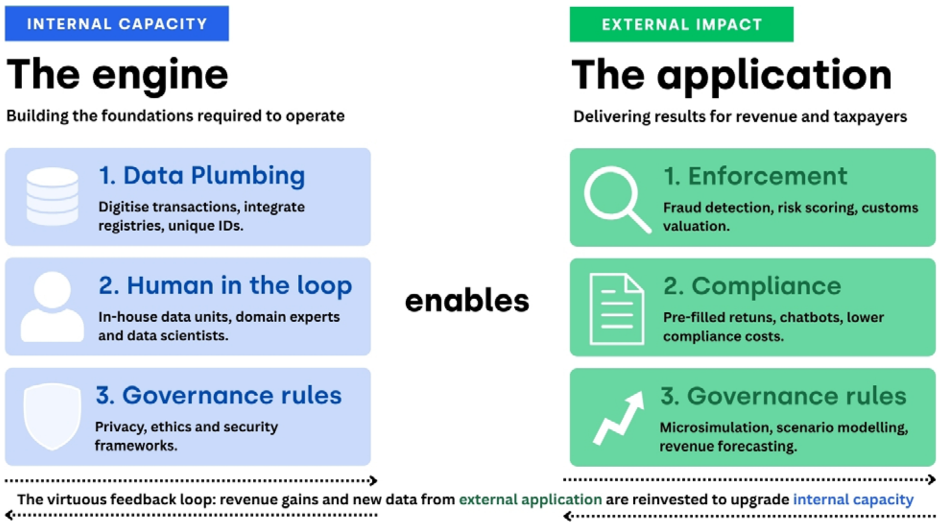

Operational Roadmap: Moving up the Ladder

This toolkit has shown how AI and complementary digital tools can significantly improve revenue recovery by strengthening compliance, sharpening risk selection, and making operations more efficient. It mapped where these tools are most promising, from basic rule-based validation to advanced machine learning and vision systems, and introduced a digital maturity framework that helps administrations diagnose readiness and sequence investments. The paper also set out practical design principles for pilots, data governance and ethical safeguards, and highlighted the organisational changes required to embed analytics into everyday work. Implementing these recommendations will require investment in data foundations, changes to institutional processes, targeted upskilling of staff, and a commitment to continuous improvement backed by clear metrics and accountability.

Tax administrations should treat AI and digital tools as a continuous transformation, not a one-off project, by building reliable data foundations, a small in-house analytics capability, and clear governance before scaling advanced models. Begin with practical digitisation and data integration at Level 1, use targeted pilots that are linked to measurable ROI and operational workflows, and require simple readiness checks and privacy reviews before any new automation. Adopt a learning strategy that iterates quickly, reinvests efficiency gains into higher complexity use cases, and reports impact routinely to maintain executive and political support. The roadmap that follows sets out specific steps across processes, data and infrastructure, analytics and technology, governance, and organisation and skills to make this progression systematic and low risk.

As a tax administration continues its digital journey, the tools it develops also need to filter through to frontline staff. Embedding AI and digital tools in core workflows and processes is vital to building a data-driven culture in the tax authority. Reducing the operational barriers between tax officials and these models is key (Davis, 1989). Increasing the amount of supporting training and information about how these models are trained and generate results provides a way for tax officials to challenge model outputs if they don't look correct by applying their context-specific expertise. This is the optimal model for working with AI – letting the tools do the hard work and using expert human judgment for verification.

For revenue authorities at Level 1 of the digital maturity framework, focus first on digitising transactions through EBMs and e-filing and basic data integration, not on advanced models. External research collaborations (with local universities, the IGC, World Bank/IMF/UN bodies, etc.) are particularly important at this level to experiment with machine learning models on real data, with a clear path to operationalisation. Tax administrations should prioritise building a small risk analytics team within the administration that combines domain experts with data scientists (this could form the foundation of a new AI team). Combining sources of data and starting with risk-scoring is a highly effective first step.

Transitioning from Level 1 to Level 2 requires a tax administration to move from paper-based siloes and ad-hoc processes to digitisation of basic registries, establishing unique identifiers for taxpayers, and instituting simple data quality controls. The following is a structured roadmap – across the six dimensions identified in the Digital Maturity Framework in the Appendix – to guide this progression:

Dimensions of digital maturity	Moving from L1 to L2
Processes	• Digitise priority end-to-end processes – select 1 or 2 high volume processes such as registration and filing, and move them to fully digital submission and acknowledgement • Replace repeated manual copying with light automation such as spreadsheet macros or simple workflow queues, so staff spend less time on routine tasks • Establish measurable process targets and short feedback loops – define 3 basic KPIs such as % of filings received digitally, average time to acknowledge registration, and % of returns with validated TINs
Data and infrastructure	• Make the unique taxpayer identifier the single source of truth – create a mapping table that links legacy identifiers to the unique TIN and run automated matching routines to merge records; prioritise cleaning of core registries (e.g., TIN registry, payment and filing logs, etc.) • Put a basic consolidated data store in place – consolidate cleaned extracts from registration, filing, and payments into a simple database that is refreshed on a scheduled basis; build basic Export-Transform-Load (ETL) jobs that apply schema validation and flag mismatches for manual review • Improve basic data quality controls and backups – implement routine data profiling to track completeness and common format errors; ensure regular backups and access controls before expanding digital services
Analytics and technology	• Move from ad-hoc spreadsheets to repeatable reports and simple dashboards – implement a small dashboard tracking core KPIs and common error types for operational leads • Pilot light automation that reduces manual work and improves data readiness – trial one virtual agent for high volume, low complexity queries using canned responses or scripted tree logic and integrate it with the call centre ticket log; use simple reconciliation scripts that automatically suggest matches between payments and filings for staff review • Keep the technology stack simple and maintainable – prefer low code or managed services and avoid bespoke systems that are hard to operate with small teams

Dimensions of digital maturity	Moving from L1 to L2
Governance and strategy	• Create core governance roles and a minimum digital strategy – appoint a senior sponsor for digital transformation and name data stewards for each priority registry; produce a strategy document that states objectives, priority processes for digitisation and an initial budget estimate • Put in place basic privacy and security safeguards – adopt simple access control rules, logging of system access and encryption for backups; agree a short policy for when administrative data may be used for analysis and who approves ad-hoc reports • Require a minimal readiness checklist before any automation pilot proceeds
Organisation and skills	• Build small, focused capability and on-the-job training – staff a small digital delivery team combining one IT lead and one data lead plus secondees from operations to ensure solutions fit practice; provide practical training on spreadsheet best practice • Create clear responsibilities and quick escalation paths – define who fixes data issues, who authorises a change to an intake form, who signs off a pilot result • Encourage quick wins and show impact – run short, time-bounded pilots that produce measurable time savings or error reductions and use results to build buy-in for the next investments into L2 capabilities

Transitioning from Level 2 to Level 3 requires a tax administration to move from isolated pilots and basic digitisation into fully defined, data-driven workflows.

Dimensions of digital maturity	Moving from L2 to L3
Processes	• Document and standardise – expand standard operation protocols beyond a few priority processes to cover all major functions (registration, filing, audit collections) • Formalise process maps showing data hand-offs, decision points, and exceptions • Automate routine steps – identify repeatable tasks (e.g., data entry, routine checks) and replace manual data hand-offs with automated workflows or RPA bots • Metrics and monitoring are crucial – define process KPIs (e.g., average filing turnaround, audit cycle time) and build dashboards to track them in near-real time.

Dimensions of digital maturity	Moving from L2 to L3
Data and infrastructure	• Integrate data siloes – consolidate taxpayers' returns' and customs' registries into a single data platform/warehouse. • Enhance data quality by rolling out data-profiling routines to measure completeness, accuracy, and timeliness on a daily or weekly basis. Establish remediation workflows, for example, by auto-notifying data stewards when a completeness metric falls below 80%. • Build or extend APIs so that analytic models can pull fresh data on demand, eliminating manual extracts.
Analytics and technology	• Move from descriptive to prescriptive – expand beyond dashboards by moving 1-2 ML pilots (e.g., chatbots, outlier detection) into frontline services • Introduce predictive models for risk-scoring (e.g., audit selection, refund fraud checks) and integrate these scores into case-management systems • Operationalise ML by setting up an ML operations pipeline with version control for datasets and models, automated retraining triggers based on data drift, and A/B testing frameworks • Put the outputs of models directly in officers' and agents' work interfaces instead of as standalone reports – this is much more likely to be integrated into the daily workflow of staff with less friction.
Governance and strategy	• Formalise AI governance by moving from an ad-hoc steering body to a standing AI/data governance board with clear charters, meeting schedules, and decision rights. • Adopt formal policies for model validation, ethical review, and privacy impact assessments for every new AI application • Allocate dedicated budget lines for analytics platform maintenance, data engineering and pilot scale-up rather than as one-off projects
Organisation and skills	• Transition the analytics team into a fully staffed Data Science Unit with distinct roles (Data Scientist; Data Quality Specialist; AI Product Manager; Data Engineer, IT and Cloud Administrator; and an Ethics and Compliance Lead) • Upskill core staff by rolling out mandatory "data literacy" and "model interpretation" training for frontline officers so they can trust and act on ML outputs • Define clear advancement tracks and incentives for analytics talent to create career pathways

The aim of increasing digital maturity is for the tax administration to have a greater impact on its stated objectives with technology. Depending on whether the tax authority wants to increase voluntary compliance or reduce tax evasion and fraud, the key question is how the technology is evaluated and what its impact is. Depending on the objective in question, the key metrics to measure will vary significantly. Having a clear sense in advance of what a specific tool or initiative is supposed to achieve helps the tax authority build monitoring and evaluation directly into the pilot process, allowing the authority to identify ineffective approaches early with minimal risk. This approach helps quickly identify what works and what doesn't and creates incentives for honest feedback and learning by doing.

References

Algorithmic Transparency Recording Standard template. (2025, May 8). GOV.UK. **https://www.gov.uk/government/publications/algorithmic-transparency-template**

Allingham, M. G., & Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of Public Economics*, 1(3–4), 323–338. **https://doi.org/10.1016/0047-2727(72)90010-2**

Almunia, M., Hjort, J., Knebelmann, J., & Tian, L. (2024). Strategic or Confused Firms? Evidence from "Missing" Transactions in Uganda. *Review of Economics and Statistics*, 106(1), 256–265. **https://doi.org/10.1162/rest_a_01180**

An overview of digital transformation at Receita Federal do Brasil. (n.d.). Inter-American Center of Tax Administrations. Retrieved 19 September 2025, from **https://www.ciat.org/an-overview-of-digital-transformation-at-receita-federal-do-brasil/?lang=en&utm_source=chatgpt.com**

ATO AI transparency statement. (n.d.). [Content]. Australian Tax Office; scheme=AGLSTERMS.AglsAgent; corporateName=Australian Taxation Office; address=GPO Box 9990 Canberra ACT 2600; contact=13 28 65. Retrieved 13 November 2025, from **https://www.ato.gov.au/about-ato/commitments-and-reporting/information-and-privacy/ato-ai-transparency-statement**

Bachas, P., & Jensen, A. (2023). Administrative data and methodologies for tax policy. *International Growth Centre.*

Bellon, M., Khalid, S., Chang, J., Villena, P., Paliza, J. C., & Dabla-Norris, E. (2022). Digitalization and Tax Compliance Spillovers: Evidence from a VAT e-Invoicing Reform in Peru. *IMF Working Papers*, 2022(057), 1. **https://doi.org/10.5089/9798400204876.001**

Carrillo, P., Donaldson, D., Pomeranz, D., & Singhal, M. (2023). Ghosting the Tax Authority: Fake Firms and Tax Fraud in Ecuador. *American Economic Review: Insights*, 5(4), 427–444. **https://doi.org/10.1257/aeri.20220321**

Chitty-Venkata, K. T., Mittal, S., Emani, M., Vishwanath, V., & Somani, A. K. (2023). *A Survey of Techniques for Optimizing Transformer Inference* (No. arXiv:2307.07982). arXiv. **https://doi.org/10.48550/arXiv.2307.07982**

Collin, M., Di Maro, V., Evans, D. K., & Manang, F. (2025). Property Tax Compliance in Tanzania: Can Nudges Help? E*conomic Development and Cultural Change*, 73(4), 2063–2103. **https://doi.org/10.1086/734186**

Coolidge, J. (2010, March 28). *Tax Compliance Cost Surveys: Using Data to Design Targeted Reforms.* SSRN. **https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1579462**

Davies, R. (n.d.). *Using fixed cargo scanners to combat tax evasion in Uganda.* Retrieved 19 September 2025, from https://www.theigc.org/sites/default/files/2025-02/Davies%20et%20al%20Policy%20brief%20May%202024.pdf

Davies, R. B., McNabb, K., & Palansky, M. (2024). You Can't Tax What You Can't See: Using Fixed Cargo Scanners to Combat Tax Evasion. *International Growth Centre.*

Dillon, B., Jang, J., Mulenga, J., & Siwale, T. (2023). *Incentivising small firms to boost VAT revenue: Evidence from Zambia* (No. ZMB-23011). International Growth Centre. https://www.theigc.org/sites/default/files/2025-02/Dillon%20et%20al.%20Policy%20Brief%20October%202023.pdf

Dillon, B., & Siwale, T. (2022). *A consumer incentive scheme experiment: Pilot study.*

Duarte, F. (2025, August 18). Number of ChatGPT Users (July 2025). *Exploding Topics*. https://explodingtopics.com/blog/chatgpt-users

Eissa, N., & Zeitlin, A. (2014). Incidence and Impact of Electronic Billing Machines for VAT in Rwanda. *International Growth Centre*. https://www.salesdatacontroller.com/wp-content/uploads/2015/09/Eissa-Incidence-and-Impact-Rwanda-VAT.pdf

Finan, F., Dal Bo, E., Schechter, L., & Duarte, R. (2025, February 1). *Countering customs fraud and corruption with machine learning.* International Growth Centre. https://www.theigc.org/collections/countering-customs-fraud-and-corruption-machine-learning

GovTech Maturity Index: 2022 Update—Trends in Public Sector Digital Transformation (Equitable Growth, Finance and Institutions Insight). (2022). https://openknowledge.worldbank.org/server/api/core/bitstreams/5e157ee3-e97a-5e42-bfc0-f1416f3de4de/content

Guidance on AI and data protection. (2025, September 9). ICO. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/

HMRC. (2025, June 26). *HMRC: Ask HMRC online.* GOV.UK. https://www.gov.uk/algorithmic-transparency-records/hmrc-ask-hmrc-online

HMRC: Using technology and AI to close the tax gap. (2024, June 24). Finsbury Robinson. https://www.finsburyrobinson.co.uk/articles/hmrc-using-technology-and-ai-to-close-the-tax-gap?utm_source=chatgpt.com

How we use data and analytics. (n.d.). [Content]. Australian Tax Office; scheme=AGLSTERMS.AglsAgent; corporateName=Australian Taxation Office; address=GPO Box 9990 Canberra ACT 2600; contact=13 28 65. Retrieved 19 September 2025, from https://www.ato.gov.au/about-ato/commitments-and-reporting/information-and-privacy/data-and-analytics/how-we-use-data-and-analytics?utm_source=chatgpt.com

Hur, H., & Abner, G. (2024). What makes public employees want to leave their job? A meta-analysis of turnover intention predictors among public sector employees. *Public Administration Review*, 84(1), 115–142. **https://doi.org/10.1111/puar.13601**

Jensen, A., Brockmeyer, A., & Gadenne, L. (2024). Taxation. *VoxDev.* **https://voxdev.org/voxdevlit/taxation**

Junquera-Varela, R. F., & Lucas-Mas, C. Ó. (2024). Revenue Administration Handbook. *World Bank.* https://doi.org/10.1596/978-1-4648-2053-3

Karten, S., Li, W., Ding, Z., Kleiner, S., Bai, Y., & Jin, C. (2025). *LLM Economist: Large Population Models and Mechanism Design in Multi-Agent Generative Simulacra* (Version 1). arXiv. **https://doi.org/10.48550/ARXIV.2507.15815**

Kleven, H. J., Knudsen, M. B., Kreiner, C. T., Pedersen, S., & Saez, E. (2011). Unwilling or Unable to Cheat? Evidence From a Tax Audit Experiment in Denmark. *Econometrica*, 79(3), 651–692. **https://doi.org/10.3982/ECTA9113**

Mascagni, G., Mengistu, A. T., & Woldeyes, F. B. (2021). Can ICTs increase tax compliance? Evidence on taxpayer responses to technological innovation in Ethiopia. *Journal of Economic Behavior & Organization*, 189, 172–193. **https://doi.org/10.1016/j.jebo.2021.06.007**

McKinsey & Company. (2023). Economic potential of generative AI. McKinsey & Company. **https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#/**

Mittal, S., Reich, O., & Mahajan, A. (2018). Who is Bogus?: Using One-Sided Labels to Identify Fraudulent Firms from Tax Returns. Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies, 1–11. **https://doi.org/10.1145/3209811.3209824**

OECD. (n.d.). Inventory of Tax Technology Initiatives [Data set]. OECD Data Explorer. Retrieved 13 November 2025, from **https://data-explorer.oecd.org**

OECD. (2020). *Tax Administration* 3.0: The Digital Transformation of Tax Administration. OECD. **https://doi.org/10.1787/ca274cc5-en**

OECD. (2025). *Tax Administration Digitalisation and Digital Transformation Initiatives.* OECD Publishing. **https://doi.org/10.1787/c076d776-en**

OECD Legal Instruments. (n.d.). Retrieved 13 November 2025, from **https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449**

Okwany, K. (2025, March 6). Kenya and Uganda tackle the challenge of exchanging images of scanned cargo. *WCO News.* **https://mag.wcoomd.org/magazine/wco-news-106-issue-1-2025/kenya-uganda-exchanging-images-scanned-cargo/**

RRA partners with AIMS to enhance performance and service quality through mathematical sciences. (2024, June 11). Rwanda Revenue Authority. **https://www.rra.gov.rw/en/details?tx_news_pi1%5Baction%5D=detail&tx_news_pi1%5Bcontroller%5D=News&tx_news_pi1%5Bnews%5D=2135&cHash=f3b4bb6e5451c72820deef42033dfde1**

Siwale, T., Shauri, A., & Mittra, S. (2025, May 5). Leveraging electronic fiscal devices in Tanzania. *International Growth Centre.* **https://www.theigc.org/blogs/leveraging-electronic-fiscal-devices-tanzania**

Steenbergen, V. (2017). Reaping the benefits of Electronic Billing Machines: Using data-driven tools to improve VAT compliance. *International Growth Centre.*

Stone, H. S. (1971). *Introduction to Computer Organization and Data Structures.* McGraw-Hill, Inc.

Tax authorities adopt AI for tax fraud and efficiencies—Vatcalc.com. (2025, February 8). **https://www.vatcalc.com/artificial-intelligence/tax-authorities-adopt-ai-for-tax-fraud-and-efficiencies/**

The AI Act, 2021/0106 (COD). Retrieved 13 November 2025, from **https://artificialintelligenceact.eu/wp-content/uploads/2021/08/The-AI-Act.pdf**

Trendall, S. (2023, July 24). HMRC reveals £5m cost savings from chatbot use during pandemic. *PublicTechnology.* **https://www.publictechnology.net/2023/07/24/economics-and-finance/hmrc-reveals-5m-cost-savings-from-chatbot-use-during-pandemic/**

UK Government. (n.d.). *The Government Data Quality Framework.* GOV.UK. Retrieved 19 September 2025, from **https://www.gov.uk/government/publications/the-government-data-quality-framework**

Waseem, M. (2025). When Enforcement Goes Digital: Balancing Revenue, Costs, and Welfare in Modern VAT Systems. *International Growth Centre.* **Unpublished**

Wasi, A. T., Faisal, W., Islam, M. R., & Bappy, M. M. (2024). *Exploring Possibilities of AI-Powered Legal Assistance in Bangladesh through Large Language Modeling* (Version 1). arXiv. **https://doi.org/10.48550/ARXIV.2410.17210**

World Bank Group (Ed.). (2021). *Data for better lives.* World Bank.

Appendix

IGC-CATA Questionnaire Methodology

This toolkit draws on primary data collected by the International Growth Centre (IGC) through a structured questionnaire distributed to all **46 member administrations** of the Commonwealth Association of Tax Administrators (CATA) ahead of its 45th Annual Technical Conference. The survey sought to generate comparable insights into how revenue authorities are investing in technology and data capacity, integrating digital tools and emerging technologies, benchmarking performance, and managing governance challenges.

The questionnaire covered five thematic pillars:

1. Current technology adoption
2. Capacity and maturity
3. Challenges and enablers
4. Benchmarking and collaboration
5. Governance and decision-making

The instrument consisted mainly of closed-ended questions (multiple-choice, categorical, and scaled responses) to ensure comparability across jurisdictions, supplemented by a small number of open-ended items for qualitative examples and context. This design balances analytical rigour with accessibility for respondents, aligning with international best practice in tax administration research.

Of the 46 administrations surveyed, **10 provided complete responses and a further 25 submitted partial responses**. Collectively, these inputs reflect a broad diversity of institutional contexts and digital maturity levels across the Commonwealth, offering a valuable snapshot of technology integration within tax administration.

The survey's design drew methodological inspiration from the OECD's Inventory of Tax Technology Initiatives (ITTI), which documents digitalisation efforts across more than 100 tax administrations worldwide. Like ITTI, the IGC questionnaire uses structured, low-burden, self-reported data to identify emerging trends and promising practices. However, while ITTI offers a global overview anchored in the OECD's *Tax Administration 3.0 framework*, the IGC exercise focuses more narrowly on technology investment, institutional capacity, and governance within the Commonwealth context. Its purpose is to generate actionable insights for peer learning and targeted collaboration among CATA members.

As with all self-reported surveys, a number of caveats apply. Variations in institutional definitions, digital-maturity levels, and reporting capacity introduce some differences in interpretation. The response rate – 10 complete and 25 partial responses – means results are indicative rather than representative. The cross-sectional nature of the data provides a snapshot in time rather than a longitudinal perspective. These limitations are consistent with early rounds of comparable exercises such as the OECD's ITTI survey.

Despite these constraints, the findings yield valuable insights into the readiness, challenges, and priorities of tax administrations pursuing digital transformation. The data have been collated thematically, forming the foundation for the recommendations presented in this toolkit.

Digital Maturity Checklist

The framework is based on the World Bank's Revenue Administration Handbook (Junquera-Varela & Lucas-Mas, 2024).

- Tally your ticks in each column. Identify your predominant level – this is your current "Level."
- Focus next steps on bridging the gap to the adjacent column. For example, if you're on Level 2 on most dimensions, prioritise: i) formalising data-quality routines (to progress in Data and Infrastructure); and ii) launching 2-3 new ML pilots with clear ROI (to move Analytics to Level 3)
- Reassess quarterly or after each major pilot to track progress.

Dimension	Level 1 (Initial)	Level 2 (Controlled)	Level 3 (Efficient)	Level 4 (Optimised)
Processes	All core processes are ad hoc or paper-based.	A handful of high-priority processes are documented and run consistently.	Standard operating procedures exist for most functions; process metrics (e.g. turnaround times) are tracked.	Processes are continuously improved via feedback loops; exception-handling is automated.
Data & Infrastructure	Most records are on paper or in siloed legacy systems; no data-quality monitoring.	Key registries (taxpayer, returns, customs) are digitised; basic data-quality checks (e.g. completeness rates) are in place.	Integrated data warehouse or platform; regular profiling (completeness, accuracy) and remediation workflows.	Near–real-time data flows feed ML pipelines; data-quality KPIs (e.g. > 95% completeness) are SLA-monitored.
Analytics & Technology	No analytics beyond static reports; no pilots.	Descriptive dashboards available; 1–2 ML pilots (e.g. outlier detection, chatbots) have run successfully.	Predictive models (risk-scores, pre-fill returns) in production; API-driven insights embedded in workflows.	Prescriptive and graph-ML applications are fully integrated (e.g. automated case-selection, real-time customs decisions); continuous retraining is in place.

Dimension	Level 1 (Initial)	Level 2 (Controlled)	Level 3 (Efficient)	Level 4 (Optimised)
Governance & Strategy	No formal AI or data governance; no written strategy.	An AI/data steering committee meets occasionally; a high-level "digital" strategy exists but with no dedicated funding.	AI/data strategy approved, with allocated budget; policies on privacy, ethics, procurement and change management are documented and enforced.	Governance is embedded at all levels (data-privacy by design, ethical review panels, vendor-agnostic tool accreditation, SLA-based performance monitoring).
Organisation & Skills	No in-house data or analytics expertise; skills entirely outsourced or absent.	A small analytics team exists; occasional staff training on basic data tools (Excel, BI); no formal career path.	Dedicated data-science unit with defined roles (data engineer, ML engineer, product manager); mandatory AI/data literacy training for frontline staff.	Enterprise-wide "citizen data" programmes; centre of excellence runs incubator, hosts hackathons and continuous certification for advanced AI skills.
Key Metrics & Diagnostics	Digital-filing rate < 20% No ML projects Data completeness < 60% 0% staff with AI training	Digital-filing rate 20–50% 1–2 pilots completed Data completeness 60–80% < 10% staff with basic analytics training	Digital-filing rate 50–80% 3–5 ML use-cases in production Data completeness 80–95% 10–30% staff certified in data/AI	Digital-filing rate > 80% > 5 complex AI applications live Data completeness > 95% > 30% staff in advanced AI roles

The key metrics and diagnostics provide the opportunity to measure progress against the strategic objectives of the wider digitalisation drive, allowing tax authorities to focus limited institutional capacity on measures that are quantifiable and that help directly build capabilities for future levels of digital maturity. More advanced levels of digital maturity mean more advanced tools to combat evasion and encourage compliance.

The Digital Maturity Framework

Having a framework allows tax administrations to diagnose areas for improvement, prioritise dimensions for progress, and sequence the introduction of AI and digital technology into tax administrators' workflows.

Dimension / Level	Level 1 — Initial (red)	Level 2 — Controlled (amber)	Level 3 — Efficient (green)	Level 4 — Optimised (blue)
Processes	Mostly manual workflows with paper-based registration and filing	Some functions digitised (e.g., registration, e-filing)	Most workflows digitised and automated (VAT scoring, targeting)	Real-time optimisation; dynamic interdiction and assignment
Data	Fragmented records; inconsistent taxpayer identifiers; limited cleaning	Key registries digitised; basic cleaning underway	Integrated warehouse; 360° taxpayer view; profiling	High-quality, real-time streams; graph and unstructured data
Key data types	Paper files; spreadsheets; partial digital extracts; basic payment or filing logs	TIN registry; e-filing logs; payment gateway; call/ chat logs	Merged registry+third-party data; invoice-level customs; audit logs; text corpora	POS/IoT/mobile money streams; graph databases; images/satellite
Technology	Standalone systems with limited interoperability; basic office software	Day-to-day IT; some dashboards	APIs, automation, models in production	Cloud-native ML infra, graph analytics, prescriptive AI
People	Limited digital skills; basic IT support; minimal analytics exposure	Small IT/analytics team; Excel/ Power BI training	Dedicated data science unit (Python/R); data governance roles	Enterprise data literacy; ML engineers and cross-functional teams
Governance	Informal or ad hoc decision-making; no digital strategy	High-level strategy; ad-hoc coordination	Formal AI/data governance; PIAs and ethics policies	Mature enterprise AI governance; audit trails; model cards
AI use	No AI use; rules-based processing	Descriptive/rule automation; pilots (chatbots, outlier detection)	Predictive models in production (fraud, refund risk)	Full spectrum (NLP, vision, RL) with feedback loops
Key risks	Data inconsistency; manual errors; weak controls; scattered systems	Data silos; poor quality; pilot over-reliance; bias	Model bias/ errors; privacy/ compliance; operational over-dependence	Privacy (real-time); autonomous errors; cyberattacks; opacity
Pragmatic next step	Establish core digital foundations such as unique identifiers; digitisation of key registries; basic data quality routines	Small, targeted revenue recovery pilot; build ETL and data stewardship	Validate & scale predictive models; PIA and validation checklist	Enterprise rollout; third-party audits; MLOps and continuous telemetry

Maturity models such as the above help an organisation understand its current level of functional, strategic, or organisational capability, and the types of changes required to enable higher levels of maturity over time. The OECD helpfully has a list of Tax Administration and Tax Crime Maturity Models available on its website.

Below, we outline the key characteristics and risks of each level of digital maturity in our framework.

Level 1

An administration on **Level 1** of the digital maturity framework can be characterised by the following features:

- AI use: none in production. Use of automation is limited to simple rule-based checks or manual scripting rather than machine learning.
    - Level 1 cases should focus on creating reliable digital records and lightweight automation that reduces manual work and error rates.
    - For example, an e-filing validation script that checks the format and presence of the taxpayer identifier before accepting a submission can eliminate large volumes of manual corrections and speed processing. A small reconciliation macro that flags obvious payment mismatches can save staff days each week and create the reliable inputs needed for future analytics.

Risks at Level 1:

- Fragmented and inconsistent identifiers and records, for example multiple local identifiers or handwritten TINs in scanned forms ➔ prioritise a single unique taxpayer identifier, map legacy identifiers to it, and publish a short data standard for name, address and identifier fields so all intake systems use the same format.
- High manual error rates from paper processes and spreadsheets, leading to slow processing and weak audit trails ➔ digitise priority registries first, require digital submission for core forms where feasible, and introduce simple validation rules at data entry to catch common mistakes.
- Weak or absent data stewardship and ownership, causing unclear responsibility for fixing records ➔ assign a data steward for each registry, document owners and contact points, and set a simple weekly triage process for cleaning top priority issues.
- Low staff capacity on basic data handling and tooling, which can produce mistrust of automated outputs ➔ run short practical trainings in spreadsheet best practice and basic data cleaning, create quick reference guides and hold on-the-job coaching sessions.
- Poor IT security and backup processes expose records to loss or unauthorised access ➔ implement basic access controls, automatic backups and simple logging; where needed follow donor or central IT minimum security standards before digitising sensitive records.

- Premature pursuit of advanced analytics despite weak inputs, producing misleading outputs if attempted ➔ require a minimum data readiness checklist before authorising any ML pilot; limit initial automation to rules and heuristics until data quality and identifiers are satisfactory.

Level 2

An administration on **Level 2** of the digital maturity framework can be characterised by the following features:

- AI use: early pilots in descriptive or rules-based automation (e.g., chatbots, outlier detection).
- Level 2 cases focus on foundational or automated support, focusing on basic, structured tasks.
- For example, a system that correctly identifies the economic activity of a taxpayer can save months of manual labour and increase the administration's risk management capabilities. Using specialised machine learning segmentation models, it is possible to deduce the economic activity of taxpayers from financial information. This application is very attractive for tax authorities as it does not require labelling of data – the existing classification for taxpayers is used to train the ML algorithms. For customs administrations, ML algorithms can assign a probable economic activity to unclassified trade partners.

Risks at Level 2:

- Data siloes and inconsistency (separate systems for registry, e-filing, and payments cause mismatches and manual reconciliation backlogs) ➔ build simple ETL pipelines with quality checks to catch TIN mismatches at load time, then notify data stewards to resolve
- Poor data quality (missing or malformed TINs in transaction records) ➔ basic data governance, assigning data-steward roles for each registry
- Over-reliance on pilots (small ML proofs-of-concept may not scale) ➔ require every ML pilot to document its data sources, business owner and ROI targets
- Early ML pilots (e.g., outlier detection) can embed biases (e.g., flagging lower-income traders disproportionately if data is sparse) ➔ run bias checks on pilot outputs, disaggregate error rates by taxpayer segment, and hold a pilot governance review before scaling

Level 3

- An administration progressing to **Level 3** starts to target efficiency and risk identification. Moving from reactive use cases to proactively using AI and digital tools to identify risks, reduce compliance and collection costs, and improve data matching to bring the corpus of administrative data to bear for the purposes of tax enforcement.

The key defining features of a tax administration at Level 3 are as follows:

- AI use: predictive models in production (e.g., filing behaviour, refund fraud); AI embedded in core operations.

  Risks at Level 3:

  - Model bias and errors (predictive risk-scores may over-flag small businesses or under-flag complex corporate groups) ➔ model-risk governance, requiring every predictive model to pass a validation checklist (accuracy, false-positive rate, business impact) before deployment. Embed human-in-the-loop checkpoints for high-impact decisions
  - Privacy and compliance (use of third-party data can breach local data-protection laws) ➔ privacy impact assessments (PIAs) for any new data integration, with redaction or anonymisation where needed
  - Operational over-dependence (staff may blindly trust model outputs without understanding limitations) ➔ user training and explainability, embedding model explanations in the case management UI and run hands-on workshops for officers. Implement MLOps best practices: version control, automated drift detection, and periodic retraining schedules; maintain fallback manual workflows in case of model downtime.

Level 4

An administration at **Level 4** focuses on advanced analysis and optimised processes, moving from proactive use cases of AI to using AI to predict risks and optimise and streamline processes in response. The defining features for an administration at Level 4 are as follows:

- AI use includes the full spectrum of AI (NLP, vision, graph, reinforcement learning) with feedback loops to drive continuous improvement

  Risks at Level 4:

  - Privacy risks are heightened due to real-time streams which expose continuous personal or commercial data ➔ enforce data-privacy by design by minimising retention windows, encrypting streams end-to-end and regularly auditing data-access logs
  - Autonomous interdictions can inadvertently block fully compliant taxpayers, causing economic harm ➔ build threshold-based human-in-the-loop overrides
  - Richer datasets attract sophisticated cyber-attacks ➔ adopt enterprise security frameworks (e.g., ISO 27001), conduct regular pen tests on ML endpoints and sandbox model training/serving environments
  - Algorithmic opacity ➔ model documentation and auditing, with "model cards" describing each model's training data, scope and limitations. Conduct annual third-party audits.

Tax for Growth

IGC

Policymakers and tax administrators in developing countries face challenges in collecting tax. Yet tax is at the heart of state-building. Tax for Growth is an International Growth Centre (IGC) initiative that supports tax administrators and policymakers in generating effective approaches to make taxation work for development.

theigc.org/initiatives/tax-growth

Manson Yim @ Unsplash