



Online Webinar | 18<sup>th</sup> December 2025

# The ai accountability framework for tax: building systems that reason, record & reassure

Speaker: Vishnu Bagri, Partner, Singhvi Dev & Unni LLP

SINGHVI  
DEV & UNNI LLP

# Context

---

Understanding the urgency of AI governance in taxation

## AI is entering core tax workflows

- Risk assessment & stratification
- Anomaly detection
- Virtual assistants
- Document drafting
- Data extraction & classification

## Regulation is catching up

- OECD AI Accountability guidance
- EU AI Act
- Expanding digital reporting regimes

## Rising expectations from authorities & boards

- “We should get instant analysis in today’s world”
- “Show me your workings.”
- “Explain the AI-assisted position.”

**Key Insight:** Tax authorities worldwide are already using AI for fraud detection and risk scoring. Tax teams must ensure their AI outputs are **defensible, not just efficient**.

# The Tax Trust Gap

---



## AI is Powerful - But Opaque

Modern AI systems, especially large language models, operate as "black boxes" making decisions that even their creators cannot fully explain and concerns around safety & security.



## Tax Demands Transparency

- Repeatability of calculations
- Defensibility before authorities
- Complete traceability of decisions

*Professionals will now need to interrogate the system to understand why a model arrived at a conclusion.*

## The Core Question

Can you show **how** you arrived at the conclusion?

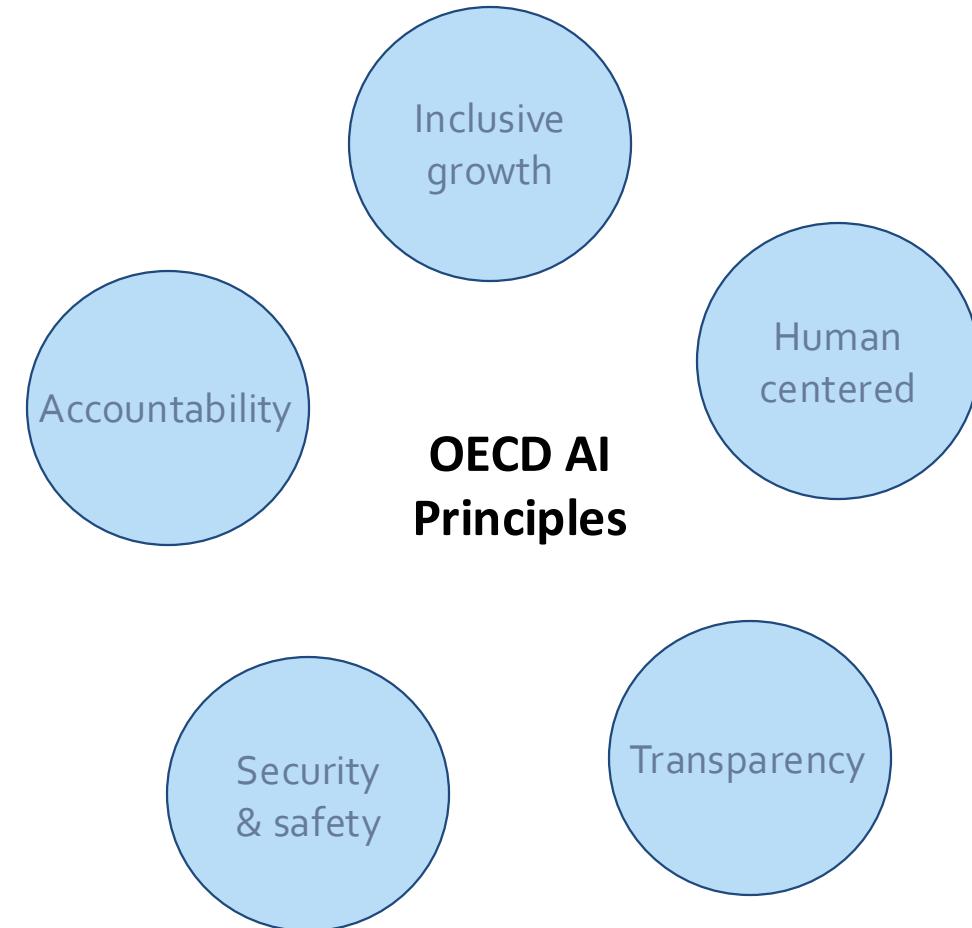
# OECD Alignment: Direction of Travel

---

How international standards shape AI governance in tax

## OECD AI Principles (2024)

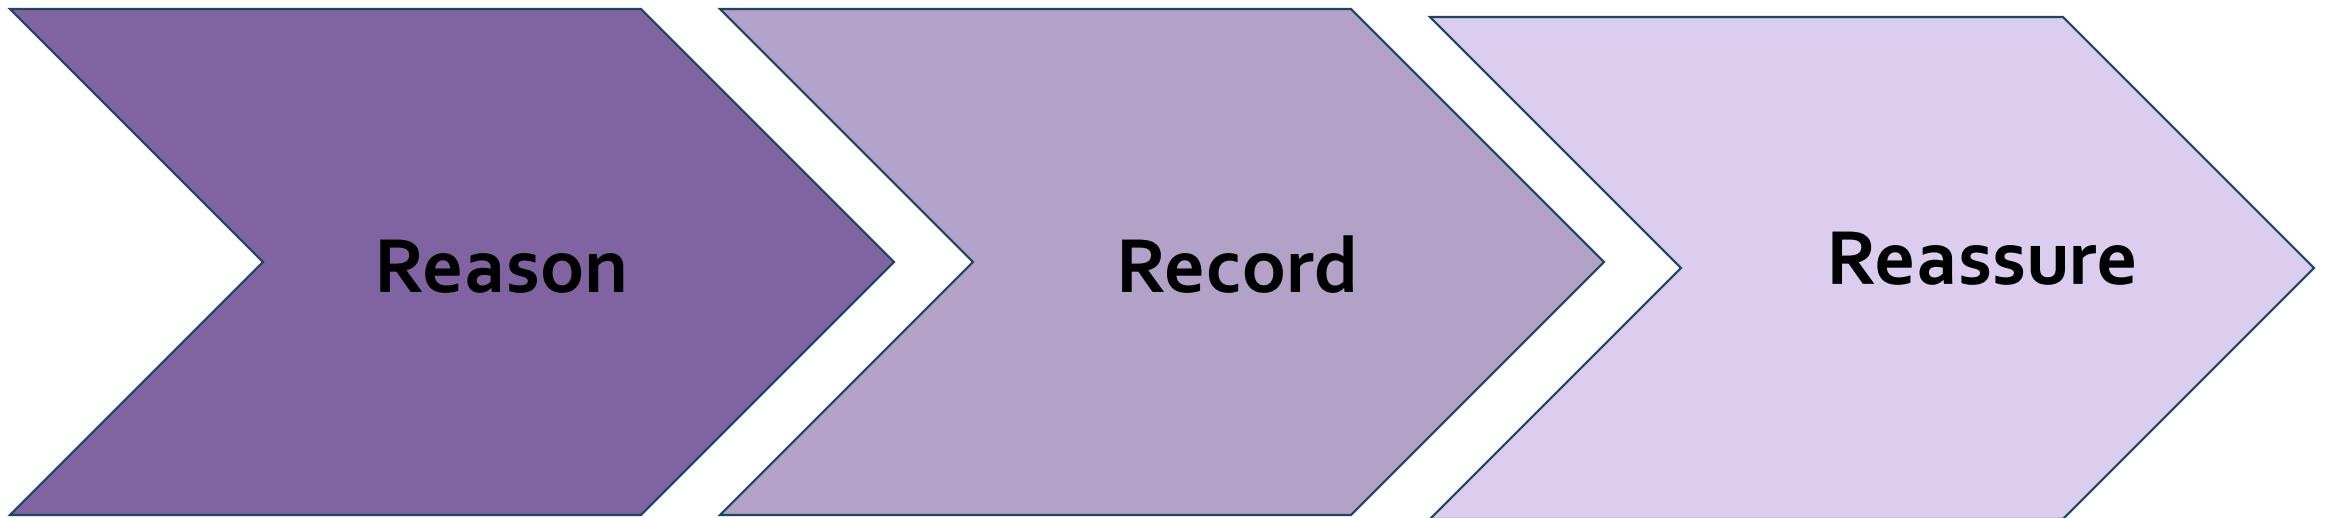
The OECD's updated AI principles emphasize that high-impact domains require explainable and reconstructible AI systems.



# The 3R Accountability Framework for AI in Tax

---

A practical model for responsible AI in tax



This framework lets tax functions adopt AI without losing control

# Pillar 1: Reason

---

## Key Requirements

### 1. Clear Rulebook

Define explicit boundaries for interpretation

### 2. Approved Sources

Specify authoritative data and legal references

### 3. Policy Alignment

Ensure AI follows organizational tax policies

### 4. Human Escalation

Route grey areas to professionals

### Why Reason Matters

AI can handle volume but struggles with nuance.

Defining rules, thresholds, and allowed sources ensures the system's behavior matches tax logic, not creative inference.

### Human Judgement needed

When AI encounters ambiguity, uncertainty, or novel scenarios, it must flag for human review rather than proceeding autonomously.

# Reason: Standardizing the workplace

---

What must be standardised is the AI-enabled workspace, not the intelligence itself

## Standardize

Input formats & templates

Mapping rules & logic

Adjustment lists

Validation criteria

## Allow Controlled Flexibility

Reviewer overrides

Interpretation notes

Exception routing

Commentary fields

## Pillar 2: Record

---

Every action the agent takes is logged: which data sources it accessed, which analysis framework it applied

### Prompts & Inputs

All queries sent to AI systems

### Model Versions

Specific versions used for each task

### Data Lineage

Source and transformation of data

### Applications

Which rules triggered and why

### Overrides

Manual adjustments with rationale

### Why Record Matters

Without logs, you cannot defend positions. Without traceability, auditors will not accept AI-enabled steps. Record is where **accountability becomes auditable**.

# Record: What a Good Audit Trail Looks Like

Creating defensible documentation



**Source Data**  
Original inputs captured

**Mapping Applied**  
Transformation rules logged

**AI Processing**  
Model decisions recorded

**Human Review**  
Overrides documented

**Final Output**  
Version history saved

| Component         | What to Capture                    | Purpose                         |
|-------------------|------------------------------------|---------------------------------|
| Source Data       | Raw files, timestamps, checksums   | Prove data integrity            |
| Rule Applications | Which rules fired, parameters used | Explain calculations            |
| Reviewer Notes    | Comments, rationale, concerns      | Document professional judgement |
| Evidence Pack     | Supporting documents, references   | Substantiate positions          |
| Version History   | All changes with timestamps        | Enable reconstruction           |

# Pillar 3: Reassure

---

Make system behavior understandable to all stakeholders



**Interpretability**  
Plain-language explanations of AI decisions



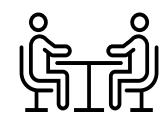
**Confidence Statements**  
Probability scores and uncertainty flags



**Role Clarity**  
Clear ownership and accountability chains

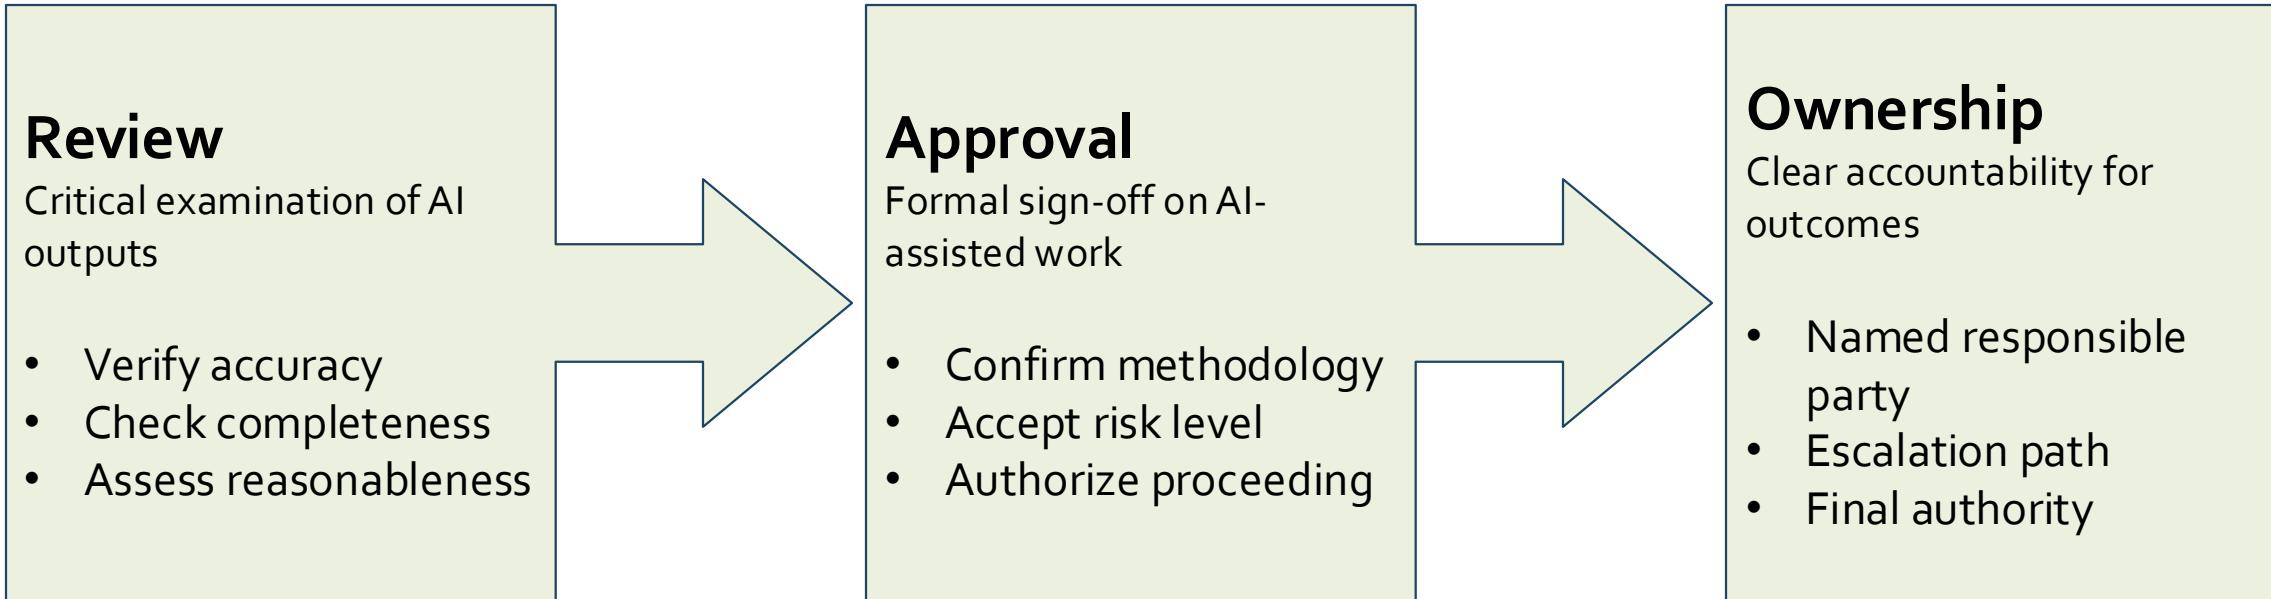


**Decision Explanations**  
Step-by-step reasoning documentation



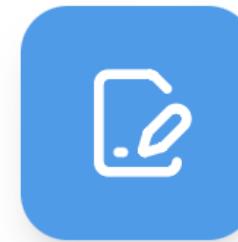
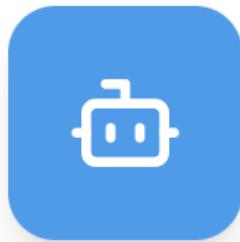
**Stakeholder Docs**  
Board and auditor-ready reports

# Accountability Nodes



Individuals responsible for reviewing, signing off, and taking ownership of outputs form the backbone of AI accountability.

# A Unified AI-Enabled Tax Workflow



## AI Preparer

Automated data processing

## AI Reviewer

Automated validation

## Human Reviewer

Professional oversight

## Sign-off

Final approval

**Fully Traceable**

Every step documented

**Repeatable**

Consistent results

**Defensible**

Audit-ready output

# AI Use Cases Across the Tax Function

---



**Corporate Tax Preparation**  
Automated computation and  
adjustments



**GST/VAT Classification**  
Smart categorization of  
transactions



**WHT Automation**  
Withholding tax rate  
determination



**Transfer Pricing Schedules**  
TP documentation and  
analysis



**Notice Response Drafting**  
AI-assisted reply preparation



**Risk Scoring**  
Anomaly detection & risk  
stratification



**Tax Control Framework**  
Continuous monitoring &  
compliance

# Deep Dive: Corporate Income Tax

---

Why CIT preparation is ideal for AI automation

## CIT Preparation

- **Rule-heavy:** Clear legislative requirements
- **Highly structured:** Standard forms and schedules
- **High volume:** Many entities, many periods
- **Year-on-year consistency:** Similar patterns repeat
- **Strong audit requirements:** Documentation essential

### OECD Perspective

The OECD identifies data preparation as one of the most automatable tasks in tax administration.

However, every extraction, mapping, and flag must be logged, otherwise this introduces blind spots.

# CIT Deep Dive: Data Gathering

## STEP 1

### How AI handles data preparation



**TB / GL Extraction**  
Automated data pull from accounting systems



**Consistency Checks**  
Validation against prior period data



**Mapping to Schedules**  
Intelligent categorization



**Missing Data Flags**  
Automated gap identification



**Prior Year Comparison**  
Variance analysis and trending

#### AI Preparer Actions

- ✓ Extract trial balance from ERP
- ✓ Map GL accounts to tax schedules
- ✓ Flag incomplete or unusual entries
- ✓ Compare with prior year patterns

#### Record Requirements

Every extraction, mapping, and flag must be logged with timestamps, source references, and the specific rules applied. This creates the foundation for an auditable workflow.

# CIT Deep Dive: Reason Layer

## STEP 2

The system applies approved tax rules across key computation areas:

**Depreciation**

- Asset categorization
- Rate application
- WDV calculations

**Disallowances**

- Personal expenses

**Interest Limits**

- Thin cap rules
- ALP requirements

**Loss Set-offs**

- Carry forward
- Set-off sequence
- Time limits

**Timing Adjustments**

- TDS compliance
- Payment dates

## Prior-Year Consistency

The system validates that rule applications are consistent with prior year treatments, flagging any changes for human review and documentation of the rationale

# CIT Deep Dive: Record Layer

## STEP 3

### Building the audit-ready evidence pack

| Log Category      | Data Captured                        | Audit Purpose         |
|-------------------|--------------------------------------|-----------------------|
| Inputs            | Source files, timestamps, checksums  | Prove data integrity  |
| Transformations   | Mapping rules, calculation steps     | Explain methodology   |
| Model Version     | AI model ID, parameters used         | Enable reconstruction |
| Rule Applications | Sections applied; thresholds checked | Justify positions     |
| Overrides         | Manual changes with rationale        | Document judgement    |
| Evidence Pack     | Supporting documents, references     | Substantiate claims   |
| Reviewer Notes    | Comments, concerns, approvals        | Record oversight      |

# CIT Deep Dive: Human Reviewer Node

## STEP 4

Where professional judgement takes over



### Compare with FS

Reconcile with financial statements



### Review Anomalies

Investigate flagged items



### Validate Rules

Confirm correct rule application



### Add Notes

Document interpretations



### Approve/Override

Accept or modify AI output



### Final Sign-off

Take accountability

The reviewer can cross-check with prior filings, detect red flags, or reconcile with financial statements

# Minimum Viable Controls (MVC)

Scaling controls with AI maturity



**For Pilots**  
Getting started

- Light logging requirements
- Human override enabled
- Basic rule set defined
- Manual review mandatory



**For Scaling**  
Growing adoption

- Role definitions established
- Rule versioning system
- Structured audit trail
- Automated validation



**For BAU**  
Full Deployment

- Integrated governance
- Evidence pack export
- Periodic testing & reviews
- Continuous monitoring

Maturity →  Sophistication

# The AI-Ready Tax Professional

Tax knowledge alone is no longer enough



## Deep Tax Knowledge

Core foundation remains essential



## Process Understanding

Know how workflows operate



## Data Literacy

Understand data structures & quality



## Interpreting AI

Evaluate algorithmic outputs



## Challenging AI

Question automated decisions

## Model Literacy

The future professional doesn't need to code.

They need the ability to:

- Question algorithmic choices
- Understand model limitations
- Identify when AI may be wrong
- Document AI-assisted reasoning

# Thank you



Rivium Quadrant 90, Capelle a/d IJssel, 2909 LC, NL

+31 10 713 8598

[welcome@e-bright.com](mailto:welcome@e-bright.com)

**SINGHVI  
DEV & UNNI LLP**