ThoughtFox

COMPLIANCE TOOLKIT

The Shadow AI Audit Checklist

A practical tool for Compliance, Risk and Governance teams in regulated organisations

April 2026 · v2

thoughtfox.ai

Shadow AI is already in your organisation

Your employees are using ChatGPT, Copilot, Gemini and a long tail of other AI tools with company data right now. Most of it is unsanctioned. Some of it is invisible to IT. None of it is waiting for your governance framework to catch up.

This checklist is a working tool, not a white paper. It is designed to be printed, shared, marked up and actioned. Work through it with your IT, security and data protection colleagues; capture what you find in the tool register on page 9; and use the "What to do next" tiers on page 11 to decide what you do on Monday morning, what you plan for the next quarter, and what you escalate to the board.

> **Before you use this — a note from ThoughtFox**
>
> This checklist is a starter guide written for learning and internal planning. It is not legal advice, regulatory guidance, or a substitute for a proper Annex III classification of the AI systems in your firm.
>
> The examples, classifications and red-flag notes are illustrative. Real decisions on EU AI Act, GDPR, sector-regulatory and contractual obligations should be taken with reference to your own legal, compliance and regulatory advisers and the specific circumstances of the systems in question.

Why compliance teams should care now

Three forces make Shadow AI a compliance issue in 2026, not a future one.

1. Article 4 has applied since 2 February 2025

The EU AI Act's AI-literacy obligation under Article 4 applied from 2 February 2025. The bulk of the Act — including most Annex III high-risk obligations — applies from 2 August 2026. You cannot evidence "sufficient AI literacy" across an organisation if you do not know which AI systems staff are using.

2. Irish sector rules raise the bar further

The Central Bank of Ireland's Consumer Protection Code 2025, and its expected role as Ireland's National Competent Authority for the AI Act in financial services, raise the bar on explainability, bias monitoring and oversight. "We didn't know that tool was

being used" is unlikely to be accepted as a defence.

3. Boards are asking

Boards and audit committees are asking the questions. They want a list of AI systems in use, the risk tier each sits in, and who owns each one. If the answer is "we're working on it", you need a credible plan — this audit is the first step.

> **You cannot govern what you cannot see. The goal of this audit is visibility first, then classification, then control.**

A few definitions before you start

A short glossary for the terms that appear most often across the rest of this checklist. Keep this handy as you work through the sections.

A few definitions before you start

Foundation model
A large, general-purpose AI model (for example GPT-4, Claude, Gemini, Llama) trained on broad data and used as the base for many different applications.

General-Purpose AI (GPAI)
Under the EU AI Act, an AI model with significant generality able to perform a wide range of tasks, regardless of how it is placed on the market. GPAI obligations for providers applied from 2 August 2025.

Deployer
A person or organisation using an AI system in the course of their activities. Most regulated firms are deployers of third-party AI tools rather than providers of their own.

Sub-processor
A third party that the AI vendor uses to process personal data on its behalf — often the underlying foundation-model provider (e.g. OpenAI sitting behind a vendor product).

Article 50
EU AI Act transparency obligations that apply to specific use cases — chatbots and AI systems interacting with people, emotion recognition, biometric categorisation, AI-generated synthetic content, and deepfakes. Not a blanket duty on every Limited-Risk system.

Annex III
The list of high-risk use cases under the EU AI Act. For financial services, the most relevant entries are 5(b) creditworthiness and credit scoring and 5(c) risk assessment and pricing for life and health insurance.

How to run this audit

A good Shadow AI audit takes four to six weeks in a mid-sized regulated firm. It is not a desk exercise. Budget time for conversations, not just spreadsheets.

Who to talk to

Shadow AI hides in the places formal governance rarely reaches. Book 15-minute conversations in the first two weeks of the audit with:

- IT and Security — the network, approved SaaS, browser extensions, firewall blocks.
- The DPO / Data Protection team — any AI-related DPIAs started or completed, and any subject-access requests that have surfaced AI-assisted processing.
- Procurement and Vendor Management — contracts signed or renewed in the last 18 months that mention AI, machine learning or automated decision-making.
- Heads of customer-facing functions (underwriting, credit, claims, onboarding, customer service) and internal functions (HR, Finance, Marketing, Legal) — both are full of quiet AI adoption.
- A sample of frontline staff — ask directly and without judgement: "What AI tools do you use to get your work done?" The answer is almost always longer than their manager's answer.

Where to look

- SaaS expense reports and credit-card statements for the last 12 months — filter for AI, GPT, Copilot, Claude, Gemini, Perplexity, Notion AI, Otter, Fireflies, Grammarly, Midjourney, Descript.
- Single-Sign-On logs and browser extension inventories — many Shadow AI tools are SSO-auth'd web apps or browser plug-ins rather than installed applications.
- The "AI", "Copilot" or "Assistant" features built into tools you already own — Microsoft 365, Google Workspace, Salesforce, HubSpot, Zoom, Slack, Notion, Atlassian. Often enabled by default.
- In-house scripts, notebooks or models built by data teams that call an AI API — these rarely make it onto any formal register.

Section A — Discovery

The first question is the simplest and the hardest: what AI tools are actually being used in your organisation, by whom, and for what?

HOW TO USE THIS SECTION (AND THE NEXT FOUR)

Tick the box where you can answer "yes" with documented evidence. Leave unticked if the answer is no, partial, or "I think so". The red-flag note shows what a concerning answer looks like. Every unticked item becomes a line on your action list.

- [ ] We have a single, current inventory of AI tools and AI-enabled features in use across the organisation.

  Red flag: No inventory exists, or the last one was compiled more than six months ago.

- [ ] The inventory distinguishes standalone AI tools (e.g. ChatGPT) from AI features embedded in existing systems (e.g. Microsoft 365 Copilot, Salesforce Einstein, Zoom AI Companion).

  Red flag: Embedded AI features are not tracked because they came with tools we already owned.

- [ ] We have asked staff directly — not just IT — which AI tools they use for work in the last 30 days.

  Red flag: The tool list comes only from IT, procurement or the intranet — not from staff themselves.

- [ ] We know which AI tools are used for decisions that affect customers (credit, claims, pricing, onboarding, servicing) and which are used in HR (CV screening, interview analysis, performance drafting).

  Red flag: Customer-facing or HR uses of AI have not been mapped to specific tools or staff.

- [ ] For each tool, we have an identified business owner — a named person accountable for its use, not just a team or function.

  Red flag: Accountability is diffuse or defaults to IT.

- [ ] Our acceptable-use policy specifically addresses generative AI, with a clear route for staff to request a new AI tool before adopting it.

  Red flag: The policy predates ChatGPT, or there is no clear request route — staff either self-serve or don't adopt.

Section B — Data going into AI tools

Every AI tool is a data-processing activity. The question is not whether data is being processed — it is which data, whose, and under what terms.

- [ ] For each AI tool in use, we know what categories of data are being entered into it by staff — confirmed by sampling real inputs, not just by policy.

  Red flag: We assume staff follow policy; we have not sampled actual prompts or inputs.

- [ ] We have confirmed in writing — from the vendor — whether customer inputs are used to train the vendor's models, and the answer is acceptable to us.

  Red flag: Nobody has read the vendor's data-use terms, or the terms are silent on model training.

- [ ] No personal data (GDPR Article 4) or special-category data (Article 9 — health, biometric, ethnicity, etc.) is being entered into consumer-grade AI tools without a lawful basis, a DPIA, and a suitable contract.

  Red flag: Staff paste customer or HR details into public AI tools to "just summarise" or "translate".

- [ ] No confidential business information — board papers, financial results before publication, M&A materials, legally privileged content — is being entered into AI tools without explicit sign-off.

  Red flag: Executives use AI to draft board materials on their personal accounts.

- [ ] For regulated activities (client money, suitability assessments, AML investigations, claims decisions), the data flowing into AI tools is documented and auditable.

  Red flag: Regulated workflows quietly route data through an AI tool with no audit log.

- [ ] We have a process for deleting data from AI tools — including conversation history — when a staff member leaves or a customer relationship ends.

  Red flag: We do not know how to delete conversation history, or the vendor does not support it.

- [ ] Staff have been trained in the last 12 months on what data is and is not safe to enter into which tools, with worked examples.

  Red flag: The last training was generic "AI awareness" with no examples, or was never delivered.

Section C — Contracts and DPAs

Every AI tool processing company or customer data needs a contract you can point to, a Data Processing Agreement that covers AI specifically, and a documented risk assessment.

- [ ] For every AI tool in use, we hold a current written contract in the organisation's name — not a personal account or a credit-card subscription expensed back.

  Red flag: AI tools are expensed personally; there is no corporate contract to point to.

- [ ] A GDPR-compliant Data Processing Agreement is in place with every AI vendor that processes personal data on our behalf, and it addresses AI-specific risks (model training on our data, sub-processors, data location, deletion).

  Red flag: DPAs are missing, out of date, or the vendor's generic template is silent on AI-specific processing.

- [ ] We have confirmed the vendor's data-residency commitments and, where relevant, international-transfer safeguards (SCCs, adequacy).

  Red flag: Data is being sent to models hosted outside the EU with no documented transfer mechanism.

- [ ] We have a documented list of each AI vendor's sub-processors, including the underlying foundation-model provider where different from the visible vendor (e.g. a tool that uses OpenAI or Anthropic behind the scenes).

  Red flag: The sub-processor chain is unknown or stops at the visible vendor.

- [ ] Liability, indemnity and service-level terms in each AI contract have been reviewed by legal and are proportionate to how the tool is actually used.

  Red flag: AI contracts were click-through accepted by individual staff without legal review.

- [ ] A Data Protection Impact Assessment has been completed for each tool that processes personal data, reviewed in the last 12 months, and we have a process to revisit it when the vendor materially changes model, terms or sub-processors.

  Red flag: No DPIAs exist, or DPIAs exist only for tools that predate generative AI.

Section D — Regulatory exposure

Shadow AI multiplies regulatory surface area. The same tool can trigger obligations under GDPR, the AI Act, sector regulation, and consumer-protection rules simultaneously. Map every tool against every regime it could touch.

GDPR

- [ ] Each AI use case has a documented lawful basis under GDPR Article 6, and our privacy notices have been updated to reflect AI processing of personal data.

  Red flag: Consent is assumed but not collected, or privacy notices still reference pre-AI processing only.

- [ ] Where AI is used to make or significantly influence decisions about individuals, the Article 22 automated-decision-making obligations have been assessed.

  Red flag: AI-assisted decisions (credit, claims, hiring) are treated as human decisions without analysis.

EU AI Act

- [ ] Every AI tool in the inventory has been preliminarily classified against the Act's four risk tiers (Prohibited, High, Limited, Minimal). Section E goes deeper on this.

  Red flag: No classification attempted, or classification is based on vendor marketing claims alone.

- [ ] Article 4 (AI literacy) has applied since 2 February 2025. We can evidence role-appropriate literacy for every role that interacts with AI — with depth differentiated by the decisions that role influences.

  Red flag: A single "AI training" module exists for all staff regardless of role.

- [ ] Prohibited-practice risks (Article 5) — social scoring, emotion recognition in workplace contexts, biometric categorisation, manipulative systems — have been actively screened for.

  Red flag: The team assumes no prohibited practices are in use because "we wouldn't do that".

Sector regulation

- [ ] CBI expectations — Consumer Protection Code 2025 explainability, the Individual Accountability Framework (including the Senior Executive Accountability Regime, SEAR), outsourcing guidance — have been mapped to each relevant AI use case, and the AI vendor inventory reconciles with the DORA third-party register.

  Red flag: CBI expectations are treated separately from AI governance, or the DORA register and AI inventory do not reconcile.

- [ ] Other sector obligations (Solvency II, MiFID II, CRD, NIS2) have been checked for AI-relevant requirements that apply to your firm.

  Red flag: AI obligations are assumed to sit entirely under the AI Act, with no cross-reference to prudential or sector regimes.

- [ ] Financial Services firms: we understand the interaction between the AI Act and existing sector requirements — particularly Solvency II model governance for insurers and MiFID II algorithmic-trading rules for investment firms.

  Red flag: The AI Act is being treated as a standalone regime rather than layered on top of existing prudential obligations.

A PRACTICAL TIP

Track the answer to every unticked item in a single action log with three fields: owner, due date, evidence when closed. If the audit cannot point to an evidence file for a closed item, it is not closed. Regulators — and boards — will ask to see the file.

Section E — EU AI Act classification

Annex III of the AI Act is where most financial-services and insurance AI lands. A Shadow AI audit that does not end with an initial Annex III screen is typically incomplete.

> **ANNEX III — THE SHORT VERSION**
>
> Annex III lists use cases that are High-Risk under the AI Act even where the underlying tool is a generic one. For this audience, the critical domains are 5(b) — AI systems used to evaluate creditworthiness or establish credit scores — and 5(c) — risk assessment and pricing for life and health insurance. Both fall under "access to and enjoyment of essential private services".

- [ ] We have screened every tool in the inventory against Annex III — not only those built for a regulated purpose. (ChatGPT drafting credit or insurance rationales is an Annex III question.)

  Red flag: General-purpose tools are assumed Limited Risk regardless of what they are used for.

- [ ] For any tool flagged as potentially High-Risk, we have documented the classification rationale — including reasons to rebut classification where we believe an exception applies.

  Red flag: Classifications are verbal or held only in one person's head.

- [ ] For High-Risk tools we are the deployer of, we have mapped the deployer obligations (human oversight, input-data monitoring, logging, informing affected persons). Where Article 50 applies — for chatbots, synthetic content, emotion recognition, biometric categorisation, or deepfakes — we have implemented the relevant transparency obligations.

  Red flag: Obligations are treated as the vendor's problem, or customer-facing chatbots do not declare themselves.

- [ ] General-purpose AI models in our stack (foundation models, embedded Copilots) have been assessed against the GPAI obligations that applied from 2 August 2025 — including where our own deployment creates new obligations beyond the model provider's.

  Red flag: GPAI obligations are assumed to sit entirely with the model provider.

Shortcut: if you would like help with classification, the ThoughtFox AI Act Classification Assistant at thoughtfox.ai/tools/ai-act-classifier takes one system from description to documented classification in roughly 20 minutes.

Your AI tool register

Capture what the audit finds in a single register. One row per tool. Print this page, or copy the structure into your GRC platform or a spreadsheet. Review monthly for the first quarter, then quarterly.

The examples below are illustrative — classifications depend on specific use cases, configurations and your firm's profile.

REGISTER FIELD GUIDE

Tool / feature: distinguish standalone apps from features of existing tools. Used by: named function, not "everyone". Data in: the categories that actually get entered — not the categories the policy permits. DPA: Yes / No / Check / N/A. AI Act tier: Prohibited / High / Limited / Minimal / Needs review. Action: the single next step, not a wish-list.

Tool / feature	Vendor	Used by	Data going in	DPA?	AI Act tier	Action
ChatGPT (consumer)	OpenAI	Ad hoc — multiple	Often personal / confidential	No	Needs review	Restrict or move to enterprise tier
M365 Copilot	Microsoft	Finance, HR, Legal	Internal documents, email	Yes — M365	Limited / review	Confirm Article 50 labelling where applicable
Zoom AI Companion	Zoom	All staff	Meeting audio, transcripts	Check	Limited	Default off for client calls
In-house credit model	Internal	Credit Risk	Customer personal + financial	N/A	High-Risk (Annex III 5b)	Full deployer obligations
Claims triage assistant	Third-party FS vendor	Claims Ops	Claimant data, medical notes	Yes — review AI clauses	High-Risk	Human oversight + logging
Writing assistant	Grammarly	Marketing, Sales	Drafts, sometimes customer names	Check	Minimal / Limited	Corporate licence + policy
[Add your own]						
[Add your own]						
[Add your own]						

What to do next

A full Shadow AI programme is quarters of work. Sequencing matters. Pick from the three tiers below rather than trying to do everything at once.

This week — quick wins

- Send a no-blame, anonymous 5-question staff survey asking which AI tools they use for work. Leave a free-text box. Close after a week.
- Pull SaaS expense and SSO logs for the last 12 months and filter for the AI vendor list on page 4. Reconcile against the register.
- Issue a one-page interim guidance note: what staff must not enter into consumer-grade AI tools until the audit completes. Keep it to three rules, not thirty.
- Check that Microsoft 365 Copilot, Zoom AI Companion, Google Workspace AI features, and any CRM-embedded AI are configured in line with your data-handling standards — defaults are rarely right.

This quarter — remediation

- Complete Sections A–E of this checklist for every tool on the register and close the red-flag items with a named owner and date.
- Put DPAs in place for every AI vendor processing personal data; retire tools where no DPA is achievable.
- Run an Annex III screen on every tool that touches customers or staff in a consequential way. Document the rationale either way.
- Update the acceptable-use policy to address AI specifically, with worked examples of "safe", "ask first", and "never".
- Deliver role-appropriate Article 4 AI-literacy training — Article 4 has applied since 2 February 2025, so role-appropriate evidence should already be in place. Board, compliance, frontline, and technical roles each need a different depth.
- Stand up a monthly AI governance forum — compliance, IT, DPO, procurement, one business owner — to keep the register current.

When to bring in outside help

Most compliance teams can run the discovery and data sections of this audit in-house. External support earns its keep in three places.

- When Annex III classification is genuinely contested — e.g. a general-purpose tool used in a regulated workflow — and you need a defensible, documented rationale before a regulator asks.
- When Article 4 literacy needs a CPD-accredited, role-differentiated programme rather than a one-size-fits-all course.
- When the audit surfaces organisation-wide patterns — AI being adopted ahead of governance — that need a transformation response rather than a compliance fix.

ThoughtFox's AI Act Classification Assistant is free to use at thoughtfox.ai/tools/ai-act-classifier. Start there with one system you are unsure about.

Sources and further reading

The primary references used to compile this checklist. Every claim in the preceding pages should be independently verifiable against these sources, which we recommend treating as your first stops for deeper reading.

EU AI Act

- Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence — eur-lex.europa.eu/eli/reg/2024/1689/oj
- Article 4 (AI literacy) — applicable from 2 February 2025.
- Article 5 (Prohibited AI practices) — applicable from 2 February 2025.
- Article 50 (Transparency obligations for providers and deployers of certain AI systems) — applicable from 2 August 2026.
- Annex III (High-risk AI systems), in particular 5(b) and 5(c) for financial services and insurance.
- European Commission AI Office — ec.europa.eu/commission/presscorner and digital-strategy.ec.europa.eu/en/policies/ai-office.

GDPR and Irish data protection

- Regulation (EU) 2016/679 (GDPR) — eur-lex.europa.eu/eli/reg/2016/679/oj
- European Data Protection Board guidance on AI and data protection — edpb.europa.eu
- Data Protection Commission (Ireland) guidance — dataprotection.ie

Irish financial-services regulation

- Central Bank of Ireland Consumer Protection Code 2025 — centralbank.ie
- Individual Accountability Framework (IAF) and Senior Executive Accountability Regime (SEAR), Central Bank of Ireland.
- Central Bank of Ireland Cross-Industry Guidance on Outsourcing (2021).
- Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA).

Brought to you by Marie Toft and the ThoughtFox team

This checklist was produced by ThoughtFox for compliance, risk and governance professionals in Irish and EU regulated firms. It draws on Marie Toft's Compliance Institute of Ireland webinar "Navigating the Ethical Frontier" (April 2026) and on ThoughtFox's engagements with financial-services, insurance and other regulated organisations navigating the EU AI Act.

We help organisations make AI work for their business, not just exist in it. We believe AI adoption is not a technology problem; it is a people opportunity. That is why we focus on building internal capability rather than dependency — so your teams own the knowledge and the systems long after our engagement ends.

If any of this audit raised questions

Reply to the email that sent you here, or reach Marie directly. We will always tell you straight whether ThoughtFox is the right fit for what you need next.

Marie Toft, AI Implementation Lead

marie@thoughtfox.ai

thoughtfox.ai | thoughtfox.ai/tools/ai-act-classifier

Important

This checklist is general guidance written by ThoughtFox for learning and internal planning. It is not legal advice, regulatory guidance, or a substitute for a proper Annex III classification of the AI systems in your firm.

The red-flag notes, illustrative classifications and suggested actions should not be relied on in place of advice tailored to your specific circumstances. EU AI Act, GDPR, Irish Central Bank and sector-regulatory decisions should be taken with reference to your own legal, compliance and regulatory advisers.

ThoughtFox accepts no liability for decisions taken or not taken in reliance on this checklist. The Act is a live, evolving instrument — verify dates, articles and guidance against the primary sources listed on the previous page before acting.