Syntra Automate

SYNTRA ASSURANCE

AI Governance for UK Businesses

A Practical Framework for 2026 — building an AI governance programme that satisfies the EU AI Act, UK regulatory expectations, and board-level accountability.

Sam Hawkins

Former British Army Intelligence Analyst · 15 years in defence, energy & AI

Head of AI, MV Global · Founder, Syntra Automate

PUBLISHED MAY 2026

SYNTRA ASSURANCE SERIES · NO. 01

CONTENTS

In this report

— *Executive Summary* 03

01 **The AI Governance Imperative** 04

02 **The Regulatory Landscape** 07

03 **The Syntra Governance Framework** 11

04 **Implementation Roadmap** 18

05 **Common Mistakes & How to Avoid Them** 20

06 **The Business Case for Governance** 22

— *About Syntra Automate* 24

— *Appendix — Templates & Checklists* 25

— *References & Further Reading* 28

ABOUT THIS REPORT

A practical, implementable framework for AI governance designed for UK mid-market businesses. It draws on direct experience delivering governance and assurance in defence, energy, and professional services. It is provided for information only and does not constitute legal advice.

EXECUTIVE SUMMARY

The governance gap is now a board-level risk

The EU AI Act is now law. The compliance deadline for high-risk AI systems is 2 August 2026 — less than three months away. UK businesses that serve EU customers, process EU citizens' data, or deploy AI systems classified as high-risk under the Act are firmly in scope.

The penalties for non-compliance are severe: fines of up to **€35 million or 7% of global annual turnover**, whichever is greater.

76%
OF UK BUSINESSES USING AI HAVE NO PROCESS TO MANAGE AI-RELATED RISKS

€35M
MAXIMUM EU AI ACT PENALTY, OR 7% OF GLOBAL TURNOVER

2 Aug 2026
HIGH-RISK AI COMPLIANCE DEADLINE UNDER THE ACT

Yet the reality on the ground is stark: 76% of UK businesses using AI have no process or practice in place to manage AI-related risks (UK Cyber Security Breaches Survey, 2025). Most organisations have adopted AI tools — from customer service chatbots to automated document review — without a corresponding governance framework. Shadow AI, where employees use tools like ChatGPT, Claude, and Microsoft Copilot without IT oversight, is rampant across every sector. This is not a theoretical risk. It is a live, measurable business exposure.

This whitepaper provides a practical, implementable framework for AI governance designed for UK mid-market businesses. It draws on direct experience delivering governance and assurance in defence, energy, and professional services — environments where getting this wrong has serious consequences.

WHAT YOU WILL FIND IN THIS WHITEPAPER

- A clear explanation of the regulatory landscape: EU AI Act, UK regulatory expectations, and international standards.
- The Syntra Governance Framework: a five-pillar approach to building a programme from scratch.
- A 90-day implementation roadmap, common mistakes to avoid, and template outlines for every core governance document.

KEY TAKEAWAY

The businesses that build governance now will not only avoid fines — they will be better positioned to scale AI safely, win contracts in regulated industries, and build lasting trust. The businesses that wait will be explaining to a regulator why they did not act when the obligations were clear.

01

CHAPTER
ONE

The AI Governance Imperative

AI governance is not a future problem. It is a present-day business risk being underestimated by the majority of UK organisations — and the regulatory environment in 2026 no longer allows for "we'll deal with it later."

1.1 WHY GOVERNANCE MATTERS NOW
1.2 THE REGULATORY CONVERGENCE
1.3 SHADOW AI
1.4 THE COST OF INACTION

Why governance matters now

Let me be direct: AI governance is not a future problem. It is a present-day business risk that is being underestimated by the majority of UK organisations.

In the last 18 months, we have seen AI move from an emerging technology to a core business tool. According to the UK Department for Science, Innovation and Technology's AI Activity Survey (2025), **68% of large UK businesses and 34% of SMEs have adopted at least one AI application**. Generative AI tools are now embedded in email, document creation, customer service, recruitment, finance, and legal workflows.

But adoption has outpaced governance by a significant margin. Most businesses adopted AI tools the same way they adopted email in the 1990s — organically, without policy, without risk assessment, and without board oversight. That was understandable when AI was experimental. It is no longer acceptable when AI systems are making decisions that affect customers, employees, and regulatory compliance.

The regulatory convergence

What makes 2026 different from any previous year is the convergence of multiple regulatory forces bearing down on AI simultaneously:

The EU AI Act (Regulation (EU) 2024/1689) is the world's first comprehensive AI law. It entered into force on 1 August 2024, with obligations phased in over three years. The critical deadline for high-risk AI systems is 2 August 2026. UK businesses are not exempt simply because they are outside the EU — the Act applies to any organisation whose AI system outputs are "used" within the EU (Article 2(1)(c)).

UK sector regulators have been given a mandate to apply AI-specific rules within their existing remits. The FCA, ICO, CMA, and Ofcom are all actively publishing guidance and, increasingly, taking enforcement action related to AI.

UK GDPR already imposes significant obligations on AI systems that process personal data, including Data Protection Impact Assessments, lawful basis for processing, and the right to meaningful information about automated decision-making (Articles 13–15, 22).

> *This is not a single regulation you can address in isolation. It is a regulatory ecosystem that demands a structured, cross-functional governance response.*

Shadow AI: the risk you cannot see

Of all the AI risks facing UK businesses today, shadow AI may be the most dangerous because it is the least visible. Shadow AI refers to the use of AI tools by employees without the knowledge, approval, or oversight of the IT or compliance function. This includes:

- Staff pasting confidential client data into ChatGPT to draft reports or emails.
- Recruitment teams using AI screening tools sourced independently, without procurement review.
- Finance teams using AI-powered spreadsheet add-ins to generate forecasts.
- Legal teams using AI to draft contracts without checking outputs for hallucinated clauses.
- Marketing teams generating content with unclear intellectual property implications.

11%

A 2025 Salesforce survey found **over half of employees using generative AI at work did not know how to use it effectively or safely, and had received no formal guidance from their employer**. A separate Cyberhaven study reported that 11% of data pasted into ChatGPT by employees is confidential.

The implications are serious. Shadow AI creates uncontrolled data exposure, potential GDPR breaches, intellectual property leakage, and — critically — it means your organisation is deploying AI systems you do not know about and therefore cannot govern.

Real-world consequences

This is not hypothetical. AI governance failures are already producing real consequences:

- **Mata v. Avianca (2023, US):** A $5,000 joint and several sanction was imposed on the two attorneys and their firm for submitting a brief with fabricated ChatGPT-generated citations. In 2024, multiple US and UK courts issued costs orders and sanctions against lawyers who filed AI-generated citations without verification, collectively reaching six figures.
- **Grok deepfake scandal (2025–26):** Grok AI generated realistic deepfakes of public figures without safeguards, prompting regulatory investigations across jurisdictions.
- **Samsung generative AI ban (2023):** Samsung banned internal generative AI use in May 2023 after engineers pasted proprietary semiconductor source code and internal meeting notes into ChatGPT.
- **ICO enforcement (ongoing):** Notices issued to organisations deploying AI on personal data without adequate transparency or impact assessments.

The cost of inaction versus the cost of governance

Business leaders often frame governance as a cost. That framing is incomplete. The cost of governance is measurable and manageable; the cost of inaction includes regulatory fines (up to €35M or 7% of turnover under the EU AI Act, £17.5M or 4% under UK GDPR), lost contracts, reputational damage, insurance exposure, and operational disruption.

KEY TAKEAWAY

Governance is not about stopping AI. It is about using it safely, accountably, and competitively. The convergence of the EU AI Act, UK sector regulators, and UK GDPR means AI governance is now a board-level business risk — not an IT project.

02

CHAPTER
TWO

The Regulatory Landscape

Three layers of regulation that UK businesses need to consider: the EU AI Act, UK domestic regulation, and international standards. Understanding it is essential — but it does not need to be complicated.

2.1 THE EU AI ACT

2.2 UK REGULATORY APPROACH

2.3 INTERNATIONAL STANDARDS

2.1 The EU AI Act

WHAT IT IS

The EU Artificial Intelligence Act (Regulation (EU) 2024/1689) was formally adopted on 13 June 2024 and entered into force on 1 August 2024. It is the world's first comprehensive legislative framework for AI and establishes a risk-based approach to regulating AI systems across the European Union.

WHO IT APPLIES TO

This is where many UK businesses make a critical error of assumption. The EU AI Act applies to:

- **Providers** of AI systems placed on the EU market or put into service in the EU — regardless of where the provider is established (Article 2(1)(a)).
- **Deployers** of AI systems established or located within the EU (Article 2(1)(b)).
- **Providers and deployers** in a third country (including the UK), where the **output produced by the system is used in the EU** (Article 2(1)(c)).

That third point is decisive. If your AI system produces outputs — decisions, recommendations, content, classifications — that are used by EU customers, EU employees, or within EU markets, you are in scope. This catches a far wider range of UK businesses than many realise.

THE RISK CLASSIFICATION SYSTEM

The Act classifies AI systems into four risk tiers:

EU AI ACT RISK TIERS — ARTICLE 5, 6, 50

TIER	DEFINITION & EXAMPLES	STATUS
Unacceptable	Banned outright: social scoring, real-time remote biometric ID, manipulative AI, workplace emotion recognition.	Prohibited from 2 Feb 2025
High	Permitted but heavily regulated: biometrics, critical infrastructure, education, employment, essential services, law enforcement, justice (Annex III).	Full compliance by 2 Aug 2026
Limited	Transparency duties: chatbots must disclose they are AI; synthetic media must be labelled.	Applies from 2 Aug 2026
Minimal	All other systems. No specific obligations; voluntary codes encouraged.	No mandatory duties

If your business uses AI in any high-risk area and the output reaches the EU, you likely have a high-risk system that requires compliance by **2 August 2026**.

KEY OBLIGATIONS FOR HIGH-RISK AI SYSTEMS

If you operate a high-risk AI system, the Act requires (Articles 8–15):

01 **Risk management system** (Art. 9) — a continuous, iterative process throughout the lifecycle.

02 **Data and data governance** (Art. 10) — training, validation and testing datasets must meet quality criteria including relevance, representativeness and freedom from errors.

03 **Technical documentation** (Art. 11) — drawn up before the system is placed on the market.

04 **Record-keeping and logging** (Art. 12) — automatic recording of events throughout operational life.

05 **Transparency to deployers** (Art. 13) — clear instructions for use.

06 **Human oversight** (Art. 14) — systems designed to allow effective human oversight.

07 **Accuracy, robustness, cybersecurity** (Art. 15) — appropriate levels of all three.

THE DEADLINE & THE FINES

2 August 2026: obligations for high-risk AI systems apply in full — conformity assessments, technical documentation, risk management, and EU database registration. The prohibited-practices provisions (Article 5) applied from 2 February 2025, and the European AI Office became fully operational on 2 August 2025. This is not a soft deadline.

€35M	€15M	€7.5M
OR 7% TURNOVER — PROHIBITED PRACTICE BREACHES	OR 3% TURNOVER — HIGH-RISK NON-COMPLIANCE	OR 1% TURNOVER — MISLEADING INFORMATION TO AUTHORITIES

For SMEs and startups, fines are capped at the lower of the two thresholds, but even the minimum fines represent existential risk for mid-market businesses.

2.2 UK Regulatory Approach

The UK has deliberately chosen not to replicate the EU AI Act. Instead, the UK Government's approach — set out in the AI Regulation White Paper (March 2023) and reinforced by the AI Opportunities Action Plan (January 2025) — is principles-based and sector-led.

THE FIVE CROSS-SECTOR PRINCIPLES

01 **Safety, security, and robustness** — AI should function securely and as intended.

02 **Appropriate transparency and explainability** — people should understand how AI decisions are made.

03 **Fairness** — AI should not produce discriminatory outcomes.

04 **Accountability and governance** — clear responsibility for AI system outcomes.

05 **Contestability and redress** — people should be able to challenge AI decisions that affect them.

These principles are not yet enshrined in primary legislation — but that does not make them optional. UK regulators have been explicitly directed to apply them and are doing so.

SECTOR REGULATORS IN PRACTICE

Financial Conduct Authority (FCA). The most active UK regulator on AI. Its 2024 AI Update set expectations for firms using AI in algorithmic trading, credit decisions, customer advice, fraud detection, and claims pricing. Firms must demonstrate they can explain AI-driven decisions to customers and the regulator.

Information Commissioner's Office (ICO). The primary UK regulator for AI processing personal data. DPIAs are required for high-risk processing (Art. 35); a lawful basis must be established; automated decisions with significant effects carry specific safeguards (Art. 22); data subjects must be informed. The ICO has issued enforcement notices in 2024–25 for profiling without adequate transparency.

Competition and Markets Authority (CMA). Focused on algorithmic pricing (AI-enabled tacit collusion), foundation-model market concentration, and consumer protection in AI-generated recommendations.

Ofcom. Focused on AI-generated content — deepfakes, synthetic media, and disinformation. Under the Online Safety Act 2023, Ofcom has powers to require platforms to address AI-generated harmful content.

UK GDPR AND AI

UK GDPR remains the most immediately enforceable framework for AI in the UK. Any AI system processing personal data must comply: Article 5 (processing principles), Article 22 (automated decision-making), Article 35 (DPIAs for high-risk processing), Articles 13–14 (right to be informed), Article 15 (right of access, including meaningful information about the logic of automated decisions).

2.3 International Standards

ISO/IEC 42001:2023 — Artificial Intelligence Management System. The most significant standard for business AI governance. Published December 2023, it specifies requirements for establishing, implementing, maintaining, and improving an AI management system, following the familiar Annex SL structure shared with ISO 27001 and ISO 9001.

NIST AI Risk Management Framework (AI RMF 1.0). Published January 2023, organised around four functions — Govern, Map, Measure, Manage. US-focused but widely referenced internationally.

IEEE Standards. IEEE 7000 and 7010 provide detailed technical guidance but are less widely adopted in UK business contexts than ISO 42001.

WHY ISO 42001 MATTERS MOST

For UK mid-market businesses, ISO 42001 offers the most practical path to demonstrable governance: it is internationally recognised, aligns with EU AI Act requirements, follows the same structure as ISO 27001, is certifiable through independent audit, and is increasingly referenced in procurement requirements.

KEY TAKEAWAY

UK businesses face a multi-layered regulatory environment. The EU AI Act is the most prescriptive and urgent; UK regulators and UK GDPR create overlapping obligations. Do not treat these in silos — build one governance programme that addresses them all. ISO 42001 provides the most practical integrating framework.

> *Syntra Automate's governance framework addresses the EU AI Act, UK regulatory expectations, and ISO 42001 alignment in a single, integrated programme — built to work in the real world, not governance for its own sake.*

03

CHAPTER
THREE

The Syntra Governance Framework

A practical, five-pillar governance framework you can begin implementing immediately — derived from real-world delivery in defence, energy, and professional services, adapted for the UK mid-market.

P1
INVENTORY & CLASSIFICATION

P2
POLICY & ACCOUNTABILITY

P3
RISK MANAGEMENT

P4
DOCUMENTATION

P5
MONITORING & ASSURANCE

01 PILLAR ONE

AI System Inventory & Classification

You cannot govern what you do not know about. The first step in any governance programme is a comprehensive inventory of every AI system, followed by classification according to risk.

Most organisations significantly underestimate the number of AI systems they operate. When we conduct AI audits, we typically find **three to five times more AI systems than the organisation expected**, once shadow AI, embedded AI, and third-party AI tools are included. Your inventory should capture internally developed systems, third-party AI products, generative AI tools, embedded AI within existing platforms, and AI-powered APIs and integrations.

CLASSIFICATION METHODOLOGY

Once inventoried, each system should be classified by EU AI Act risk tier, and additionally by data sensitivity (personal, special category, commercial, classified) and decision impact (informational, advisory, or decision-making).

AI SYSTEM REGISTER — FIELD STRUCTURE TEMPLATE A.1

FIELD	DESCRIPTION
System ID	Unique identifier
Provider / Owner	Who built/supplies it; named accountable business owner
Purpose	What the system does, in plain language
Data sensitivity	Personal / special category / commercial / public
EU AI Act risk tier	Unacceptable / High / Limited / Minimal
Decision impact	Informational / advisory / decision-making
EU nexus	Does the output reach EU customers or citizens?
Status	Active / under review / deprecated

Worked example. A “CV Screening Tool” (HireVue, third-party ML classification, owned by Head of HR, processing personal data, recruiting for EU-based roles) classifies as **High risk** under Annex III Category 4 (Employment), decision-making impact, EU nexus — yes.

KEY TAKEAWAY

The AI System Register is the foundation of your entire programme. Without a complete inventory, every subsequent step is built on sand. Conduct a shadow AI survey. Audit your software stack. Talk to every department. The systems you do not know about are the ones that will cause you problems.

02 PILLAR TWO

Policy & Accountability

Governance requires clear rules and clear ownership. An AI usage policy sets the rules. An accountability structure ensures someone is responsible for enforcing them.

An effective AI Usage Policy is not a lengthy legal document nobody reads. It is a clear, practical set of rules covering: **scope** (which tools are approved, prohibited, or require approval — name them); **data rules** (what data can and cannot be entered into AI systems); **output verification** (all AI outputs reviewed by a human before client-facing or decision-making use — a legal requirement under UK GDPR Art. 22 and EU AI Act Art. 14); **intellectual property**; **incident reporting**; and **consequences** with genuine teeth.

ACCOUNTABILITY STRUCTURE

- **Board level:** a named board member with overall accountability for AI risk — a business risk responsibility, not an IT one.
- **AI Governance Lead:** a named individual (CISO, Head of Compliance, or dedicated role) running the day-to-day programme and reporting to the board.
- **System Owners:** every system has a named business owner accountable for its compliance.
- **All employees:** responsible for following the policy and reporting incidents.

APPROVAL WORKFLOW FOR NEW AI SYSTEMS

01 Request — business user submits tool, purpose, data, vendor.

02 Classification — Governance Lead assigns risk tier and data sensitivity.

03 Assessment — risk assessment and DPIA for high-risk or personal-data systems.

04 Approval — approve, conditionally approve, or reject.

05 Registration — approved systems added to the register.

06 Review — defined intervals (minimum annual; more frequent for high-risk).

KEY TAKEAWAY

Policy and accountability are the backbone of governance. A clear, practical, enforced policy turns governance from aspiration into reality. But policy without ownership is meaningless. Name the accountable individuals. Make reporting lines clear. Give the policy teeth.

> *Syntra Automate delivers AI Usage Policies tailored to your organisation — not generic templates — that are practical, enforceable, and aligned with your regulatory obligations.*

03 PILLAR THREE

Risk Management

AI introduces specific categories of risk that traditional enterprise risk frameworks are not designed to capture. A dedicated AI risk management process is essential.

Risk assessment should be conducted for every system in the register, with depth proportional to risk classification. A practical assessment covers six categories: **data quality** (representativeness, currency, source documentation); **bias and fairness** (protected-characteristic impact under the Equality Act 2010, testing, ongoing monitoring); **security** (adversarial attack, access controls, encryption, penetration testing); **privacy** (lawful basis, DPIA, data subject rights, transfers); **reliability and accuracy** (error rates, consequence of error, human review, hallucination management); and **transparency** (explainability, user notification, regulatory explainability).

ASSESSING THIRD-PARTY AI SUPPLIERS

Most UK businesses deploy AI built by OpenAI, Anthropic, Google, Microsoft, and specialised vendors. This creates **supply chain risk** that must be actively managed: where data is processed and whether it trains the provider's models; security certifications (SOC 2, ISO 27001); AI governance maturity; contractual protections and indemnities; EU AI Act commitment; and business continuity. Do not take provider claims at face value — review documentation, test systems, and secure contractual protections.

PROMPT INJECTION AND ADVERSARIAL RISK

A category many businesses overlook. In December 2023 a car dealership's AI chatbot was manipulated into offering a vehicle for $1 via prompt injection. Adversarial inputs in emails can manipulate AI assistants; malicious code in repositories can influence AI coding suggestions. These are not theoretical — they are being actively exploited. Your risk assessment must include adversarial scenarios.

> KEY TAKEAWAY
>
> AI risk management is not filling in a form once. It is an ongoing, structured process covering data quality, bias, security, privacy, reliability, and transparency — for every system, including third-party tools. The supply chain is your biggest blind spot. Assess providers as rigorously as your own systems.

04 PILLAR FOUR

Documentation & Transparency

If it is not documented, it does not exist — at least as far as regulators are concerned. The EU AI Act places significant documentation requirements on high-risk systems; UK GDPR requires transparency about automated processing.

Article 11 requires providers of high-risk systems to draw up technical documentation *before* the system is placed on the market (Annex IV): general description, development methods and design choices, inputs and outputs, monitoring and control, risk management system, lifecycle changes, harmonised standards, declaration of conformity, and performance metrics. Deployers must use the system per instructions, ensure trained human oversight, monitor operation, report serious incidents, conduct a fundamental rights impact assessment (Article 27) where required, and inform affected individuals (Article 26).

DPIAS, MODEL CARDS & LOGGING

Under UK GDPR Article 35 a DPIA is required before high-risk processing — including automated decisions with significant effects, large-scale special category data, and innovative use of new technologies. Model cards and system cards are becoming standard practice; request them from third-party providers. Article 12 requires automatic logging of events for high-risk systems, with retention periods defined and documented.

HIGH-RISK DOCUMENTATION CHECKLIST (EXTRACT) TEMPLATE A.4

	DOCUMENTATION REQUIREMENT	ARTICLE
☐	General system description & intended purpose	*Annex IV(1)*
☐	Risk management system documentation	*Art. 9*
☐	Data governance measures	*Art. 10*
☐	Human oversight arrangements	*Art. 14*
☐	Accuracy, robustness & cybersecurity	*Art. 15*
☐	Logging & monitoring arrangements	*Art. 12*
☐	Data Protection Impact Assessment	*UK GDPR Art. 35*
☐	EU declaration of conformity	*Art. 47*

KEY TAKEAWAY

Documentation is not bureaucracy — it is your evidence of compliance. When a regulator asks how your system works, who approved it, and how it is monitored, your documentation is your answer. Build it from day one. Keep it current. Ensure third-party providers can give you what you need.

05 PILLAR FIVE

Monitoring & Assurance

Governance is not a project with an end date. It is an ongoing operational discipline. Point-in-time audits are necessary but insufficient. Continuous monitoring is essential.

AI systems are dynamic — they learn, update, and change, and their risks are not static. A point-in-time audit does not tell you whether models have drifted, whether new tools have been adopted, or whether policy compliance is being maintained.

CONTINUOUS MONITORING CADENCE PILLAR 5

MONITORING AREA	FREQUENCY	METHOD
AI System Register accuracy	Monthly	Procurement, IT asset, expense reconciliation
Shadow AI detection	Quarterly	Staff survey, network & SaaS review
High-risk performance	Monthly	Automated accuracy / error monitoring
Bias & fairness	Quarterly	Statistical analysis across characteristics
Security posture	Quarterly	Vulnerability & prompt-injection testing
Third-party compliance	Bi-annually	Supplier questionnaires, contract review
Regulatory horizon	Monthly	EU AI Office, ICO, FCA, CMA, Ofcom scan

Board reporting should be a clear, concise view of AI risk — not a 50-page technical report. A quarterly board report covers RAG status, AI estate overview, top five risks with trends, regulatory update, incidents, compliance status, and actions required. It should fit on two to three pages.

Incident response. Every organisation needs an AI incident response plan: Detect → Contain → Assess → Report (ICO within 72 hours for personal data breaches; EU AI Act serious incident reporting for high-risk systems) → Remediate → Learn.

KEY TAKEAWAY

Governance is an ongoing discipline, not a one-off project. Build monitoring into your operational rhythm. Report to the board in business language. Have an incident plan before you need one. Periodically get an independent pair of eyes on your programme.

> *Syntra Automate provides ongoing AI assurance — quarterly reviews, board reporting support, incident response planning, and annual independent audits — building internal capability so governance becomes self-sustaining.*

04

CHAPTER
FOUR

Implementation Roadmap

Theory is essential. Execution is everything. A practical 90-day plan for a mid-market UK business (£1M–50M revenue) with existing AI adoption but no formal governance programme.

M1	M2	M3	90+
FOUNDATION	BUILD	EMBED	CONTINUOUS IMPROVEMENT

Month 1 — Foundation (Weeks 1–4)

Objective: understand what AI you have, who is responsible, and what the rules are.

MONTH 1 ACTIVITIES — WEEKS 1–4

PHASE	KEY ACTIONS
Wk 1–2 · Inventory	Appoint AI Governance Lead; brief department heads; deploy anonymous shadow-AI survey; review IT asset, procurement and SaaS records; populate the AI System Register.
Wk 2–3 · Classification	Classify all systems by EU AI Act tier; identify EU-nexus systems; flag high-risk systems for August 2026; assess data sensitivity; prioritise for detailed assessment.
Wk 3–4 · Policy	Draft AI Usage Policy; circulate to legal, IT, HR, compliance, operations; define accountability structure; design approval workflow.

Quick win: the shadow AI survey alone typically reveals tools even IT leadership did not know were in use. This single action justifies the effort and builds organisational awareness.

Month 2 — Build (Weeks 5–8)

Objective: assess risks, build documentation, and establish controls.

MONTH 2 ACTIVITIES — WEEKS 5–8

PHASE	KEY ACTIONS
Wk 5–6 · Risk	Comprehensive risk assessments for high-risk systems; DPIAs for personal-data systems; assess third-party suppliers; document risk treatment; identify remediation gaps.
Wk 6–7 · Documentation	Build technical documentation packs; request provider documentation; establish logging requirements; create model cards; document human oversight.
Wk 7–8 · Controls	Finalise and approve the policy; design incident response plan; establish approval workflow; define monitoring metrics; set reporting cadence.

Month 3 — Embed (Weeks 9–12)

Objective: launch governance across the organisation and make it operational.

MONTH 3 ACTIVITIES WEEKS 9–12

PHASE	KEY ACTIONS
Wk 9–10 · Launch	Publish policy organisation-wide; all-staff briefing led by senior leadership; mandatory awareness training; role-specific training for high-risk operators; communicate incident reporting.
Wk 10–11 · Board	Prepare first Quarterly AI Risk Report; present status, risks, timeline, investment; secure board endorsement; agree reporting cadence.
Wk 11–12 · Sustain	First monthly AI System Register review; first monitoring cycle; capture lessons learned; plan first external assurance review; document the programme for continuity.

Post-90 days — continuous improvement

The 90-day plan gets you from zero to operational. Governance is not "done" at day 90. The ongoing rhythm: **monthly** — AI System Register review, monitoring, regulatory scan; **quarterly** — board report, shadow AI survey, policy review, bias analysis; **bi-annually** — supplier review, risk reassessment; **annually** — full programme audit, policy major review, training refresh.

KEY TAKEAWAY

Ninety days is enough to build a governance programme — enough to establish the foundation, address the highest risks, launch the policy, and begin board reporting. The critical thing is to start. Every week of delay is a week closer to the August 2026 deadline without governance in place.

> *Syntra Automate delivers 90-day governance implementations for mid-market businesses — working alongside your team, transferring knowledge, and ensuring you can sustain it independently.*

05

CHAPTER
FIVE

Common Mistakes & How to Avoid Them

The ten mistakes we see most frequently in delivering AI governance across defence, energy, and professional services. Every one of them is avoidable.

01–05
STRUCTURAL & POLICY FAILURES

06–10
OPERATIONAL & SCOPE FAILURES

Ten avoidable failures

01 **Treating governance as an IT project**

AI governance is a business risk issue. It requires board oversight, legal input, HR involvement, and operational engagement. If it lives entirely in IT, it will fail.

02 **Writing policies nobody reads or enforces**

A 40-page policy in legal language buried on the intranet is not governance. Write in plain English, keep it practical, and enforce it — a policy without consequences is a suggestion.

03 **Ignoring shadow AI**

If you have not surveyed for unsanctioned AI use, you do not know your AI estate. Shadow AI is the single biggest source of unmanaged AI risk in most organisations.

04 **Not classifying third-party AI tools**

Your CRM's lead scoring, your HR platform's candidate screening, your finance tool's forecasting — these are all AI systems, potentially high-risk. Audit the whole stack.

05 **One-off audit with no ongoing monitoring**

An audit conducted once and never repeated is like a health check you never follow up on. Systems change, risks evolve, regulations update. Build monitoring into the operational rhythm.

06 **No incident response plan for AI failures**

When — not if — an AI system produces incorrect, biased, or harmful output, you need a practised response. Write the plan, test it, train your team on it.

07 **Documentation as an afterthought**

The EU AI Act requires technical documentation before a high-risk system is placed on the market. Retrospective documentation is incomplete and a red flag for regulators.

08 **No board-level reporting**

If your board has not received a report on AI risk, governance is not on its agenda. The financial, regulatory, and reputational exposure demands board attention.

09 **Assuming UK businesses are out of scope**

The most dangerous assumption. If your AI output is used within the EU, you are in scope. Brexit did not exempt UK businesses from extraterritorial EU regulation.

10 **Waiting until the deadline**

Building a programme — assessments, documentation, training, monitoring — takes months, not weeks. Those who start now will be ready. Those who wait will not.

KEY TAKEAWAY

These mistakes are common, but not inevitable. Every one can be avoided with planning, commitment, and the right approach. If you recognise your organisation in any of the above, the time to act is now.

06

CHAPTER
SIX

The Business Case for Governance

Governance is not just about avoiding fines and satisfying regulators. There is a compelling positive business case — competitive advantage, contract eligibility, lower insurance, customer trust, and successful scaling.

6.1 COMPETITIVE ADVANTAGE
6.2 TRUST & INSURANCE
6.3 SCALING AI SUCCESSFULLY

Governance as competitive advantage

In a market where the majority of organisations have no formal AI governance, having one is a differentiator. Customers, partners, and investors increasingly ask: “How do you govern your use of AI?” Defence and government procurement now routinely includes AI governance requirements in tender evaluations; financial services firms are required by the FCA to demonstrate governance and pass it down their supply chains; professional services firms face client questions about confidentiality and data protection.

Contracts, insurance & trust

For businesses in or selling to regulated industries, AI governance is becoming a commercial necessity — ISO 42001 certification or demonstrable alignment is emerging as a procurement criterion alongside ISO 27001 and Cyber Essentials. AI-specific liability insurers are beginning to differentiate premiums by governance maturity, just as cyber insurers do for security posture.

Public concern about AI is growing rapidly. The Edelman Trust Barometer (2025) found that trust in AI companies has declined, with fewer than half of respondents in major economies expressing trust in companies developing AI technology. Organisations that can demonstrate responsible AI governance — through published policies, transparency reports, and certification — build trust that translates into customer loyalty and brand value.

Scaling AI successfully

3×

McKinsey’s 2025 State of AI report found organisations with strong senior leadership engagement in AI — including formal governance and oversight — were **three times more likely to report successful AI scaling** than those without such engagement.

The reason is straightforward. Governance forces clear thinking about what you deploy, why, and how you will manage it — preventing the chaotic, ungoverned proliferation of AI tools that leads to wasted spend, duplicated effort, and abandoned projects. It builds the organisational confidence needed to deploy AI in higher-value, higher-risk use cases. Governance does not slow you down. It speeds you up — safely.

MAXIMUM REGULATORY EXPOSURE — STATUTORY CEILINGS

REGULATION	MAXIMUM FINE
EU AI Act — prohibited practices	**€35M or 7% global turnover**
EU AI Act — high-risk non-compliance	**€15M or 3% global turnover**
UK GDPR	**£17.5M or 4% global turnover**

For a business with £20M turnover, a 7% EU AI Act fine would be £1.4M; a 3% fine £600K. The cost of a comprehensive governance programme is a small fraction of these figures.

KEY TAKEAWAY

AI governance is not just a cost of compliance — it is a commercial asset. It wins contracts, reduces insurance exposure, builds trust, and enables you to scale AI with confidence. The question is not whether you can afford to invest in governance — it is whether you can afford not to.

ABOUT SYNTRA AUTOMATE

Practitioners, not theorists

Syntra Automate is a UK-based AI governance and assurance consultancy founded by Sam Hawkins. We help mid-market businesses build AI governance programmes that satisfy regulatory requirements, protect the organisation, and enable the safe scaling of AI.

We are not an academic research firm. We are not a big consultancy selling theory. We are practitioners who have delivered governance in environments where it matters — defence, intelligence, energy, and critical national infrastructure.

OUR FOUNDER

Sam Hawkins brings a distinctive background: former British Army Intelligence Analyst, trained in intelligence analysis, risk assessment, and operational security; 15 years across defence, energy, and AI, including roles with EDF Energy, Defence Equipment & Support (DE&S), and the UK defence sector; Head of AI at MV Global, leading AI strategy, deployment, and governance; and Founder of Syntra Automate.

THE SYNTRA ASSURANCE OFFERING

SERVICE PATHWAY — SYNTRA ASSURANCE

STAGE	SCOPE & TIMELINE
1 · AI Risk Audit	Current-estate assessment, shadow AI discovery, risk classification, gap analysis vs EU AI Act / UK GDPR / ISO 42001. 2–3 weeks.
2 · Governance Framework Build	Tailored programme: register, policy, accountability, risk process, documentation, monitoring. Over 90 days.
3 · EU AI Act Compliance	High-risk documentation, conformity assessment preparation, regulatory readiness. Targeted at the August 2026 deadline.
4 · Ongoing Assurance	Quarterly reviews, board reporting, regulatory monitoring, incident support, annual independent audit.

Our approach — *"Solo-founded. AI-powered. Trusted by Defence."* — uses AI to deliver governance efficiently, builds frameworks on real-world delivery, and transfers capability so governance becomes self-sustaining.

Ready to start?

Free AI Risk Audit — understand your exposure in 60 minutes · syntraautomate.com/audit

EU AI Act Readiness Quiz — 5-minute compliance check · syntraautomate.com/quiz

Contact — syntra.automate@gmail.com · syntraautomate.com

Enterprise enquiries — sam@syntraautomate.com

APPENDIX

Templates & Checklists

Structural guides — not full templates — designed to help you understand what each governance document should contain and begin drafting. Full editable versions are available within Syntra Assurance engagements.

A.1 AI SYSTEM REGISTER

A.2 USAGE POLICY OUTLINE

A.3 RISK ASSESSMENT

A.4 DOCUMENTATION CHECKLIST

A.5 BOARD REPORT

A.1 AI System Register

AI SYSTEM REGISTER — WORKED EXAMPLE ROW TEMPLATE A.1

FIELD	VALUE
System ID	*AI-2026-014*
System name	*CV Screening Tool*
Provider	*HireVue (third-party)*
Business owner	*S. Mitchell, Head of HR*
Purpose	*Automated screening of job applications*
Data sensitivity	*Personal (potentially special category)*
EU AI Act risk tier	**High — Annex III, Cat. 4 (Employment)**
Decision impact	*Decision-making*
EU nexus	*Yes — recruits for EU-based roles*
Status	*Active · not yet reviewed*

A.2 AI Usage Policy — section headings

1. Introduction and purpose
2. Scope and applicability
3. Definitions
4. Roles and responsibilities
5. Approved AI tools and platforms
6. Prohibited AI uses
7. Data classification and handling rules
8. Human oversight and output verification
9. IP and AI-generated content
10. Procurement and approval process
11. Third-party AI supplier requirements
12. Privacy and data protection
13. Security requirements
14. Incident reporting and escalation
15. Training and awareness
16. Monitoring and compliance
17. Enforcement and consequences
18. Review schedule and version control
19. Appendix: approved tool list (quarterly)

A.3 AI Risk Assessment

AI RISK ASSESSMENT — CRITERIA MATRIX TEMPLATE A.3

ASSESSMENT CRITERIA	RATING	EVIDENCE	MITIGATION
DATA QUALITY			
Training data representative & unbiased	H / M / L	—	—
Data current & sources documented	H / M / L	—	—
BIAS & FAIRNESS			
Tested across protected characteristics	H / M / L	—	—
Ongoing bias monitoring in place	H / M / L	—	—
SECURITY & PRIVACY			
Resilient to adversarial attack	H / M / L	—	—
DPIA completed; lawful basis established	H / M / L	—	—
RELIABILITY & TRANSPARENCY			
Accuracy / consequence of error assessed	H / M / L	—	—
Outputs explainable to regulators	H / M / L	—	—

A.5 Quarterly Board AI Risk Report

BOARD REPORT — THREE-PAGE STRUCTURE TEMPLATE A.5

PAGE	CONTENT
1 · Executive summary	RAG status; headline metrics (total systems, high-risk count, compliance %); top 3 items for board attention
2 · Detailed dashboard	Estate by tier; EU AI Act readiness %, DPIA %, policy compliance %; incident summary; training completion
3 · Risks & horizon	Top 5 risks with trends; regulatory horizon; supplier status; actions/decisions required; next-quarter priorities

Full editable templates — Word documents, Excel registers, and PowerPoint board packs — are available as part of Syntra Assurance engagements, tailored to your organisation's size, sector, and regulatory obligations. Contact: syntra.automate@gmail.com

REFERENCES & FURTHER READING

Sources

01 **EU AI Act** — Regulation (EU) 2024/1689, 13 June 2024. Official Journal of the European Union.

02 **UK AI Regulation White Paper** — "A pro-innovation approach to AI regulation," DSIT, March 2023.

03 **UK AI Opportunities Action Plan** — DSIT, January 2025.

04 **ISO/IEC 42001:2023** — Information Technology — Artificial Intelligence — Management System. ISO, December 2023.

05 **NIST AI Risk Management Framework (AI RMF 1.0)** — NIST, January 2023.

06 **ICO Guidance on AI and Data Protection** — Information Commissioner's Office, updated 2025.

07 **FCA AI Update** — Financial Conduct Authority, 2024.

08 **CMA AI Foundation Models: Initial Report** — Competition and Markets Authority, 2023.

09 **UK Data Protection Act 2018** and **UK GDPR** (retained EU law).

10 **Online Safety Act 2023** — UK Parliament.

11 **UK Cyber Security Breaches Survey 2025/2026** — Department for Science, Innovation and Technology.

12 **McKinsey Global Survey: The State of AI in 2025** — findings on senior leadership engagement in AI governance and successful AI scaling.

13 **Edelman Trust Barometer 2025** — findings on declining trust in AI companies across major economies.

This whitepaper is provided for informational purposes only and does not constitute legal advice. Organisations should seek appropriate professional advice tailored to their specific circumstances. While every effort has been made to ensure accuracy, the regulatory landscape for AI is evolving rapidly and readers should verify current requirements.

Sam Hawkins

Founder, Syntra Automate · Former British Army Intelligence Analyst · Head of AI, MV Global

syntra.automate@gmail.com | syntraautomate.com ◆