SURECLOUD®

TEMPLATE

EXECUTABLE GRC

AI Governance Policy Template.

A structured, editable template for building an auditable AI governance policy — covering the seven elements regulators and auditors look for.

- Policy sections 1–9
- Risk-tier framework
- Common-gaps checklist

BEFORE YOU BEGIN

How to use this template.

This template provides the structural framework for an AI governance policy aligned with UK and EU regulatory requirements — including the EU AI Act, UK GDPR, and ISO 42001:2023.

- Complete the **Document Control** section before circulating drafts.
- Work through the nine sections in order — earlier decisions, particularly scope and risk classification, shape later ones.
- Involve your legal, compliance, and risk teams before finalising.
- Replace all placeholder text shown as [bracketed fields] with organisation-specific content.
- Treat greyed drafting guidance as instruction only — remove it from the final document.

ISO 42001

A POLICY ALONE DOES NOT SATISFY REGULATORS

Each section requires operational infrastructure behind it — an AI register, an approval workflow, audit logging, and incident reporting channels. See the companion guide at **surecloud.com/resource-hub/how-to-build-ai-governance-policy** for implementation guidance.

What this template covers

§	TOPIC	KEY OUTPUT
1	**Purpose & Scope**	Scope definition + AI system definition
2	**Risk Classification**	Three-tier risk framework
3	**Accountability**	Decision rights matrix
4	**Approval Process**	Seven-stage approval workflow
5	**AI Register**	Register structure and fields
6	**Audit & Monitoring**	Logging and review requirements
7	**Incident Response**	Incident definition and response protocol
8	**Acceptable Use**	Employee responsibilities
9	**Policy Review**	Maintenance and version control

FRONT MATTER

Document control.

Complete this section before the policy is circulated for review. It establishes ownership, version, and the related policies this document sits alongside.

Document title	AI Governance Policy
Version	[e.g. v1.0]
Policy owner	[Role — e.g. Chief Risk Officer / CISO]
Approved by	[Name, role, date]
Review date	[Date — annual minimum recommended]
Classification	[Internal / Confidential]
Related policies	[e.g. Data Protection Policy, Information Security Policy]

Revision history

VERSION	DATE	AUTHOR	SUMMARY OF CHANGES
1.0	[Date]	[Name]	Initial draft

DRAFTING GUIDANCE

Keep this table current for the life of the policy. Auditors treat a maintained revision history as evidence the policy is a live control, not a one-off artefact.

SECTION ONE

1 Purpose and scope.

1.1 Purpose

This policy governs the deployment, use, and oversight of artificial intelligence systems within [Organisation Name]. It establishes accountability structures, risk classification criteria, and approval processes to ensure AI is used responsibly, lawfully, and consistently with the organisation's values and regulatory obligations.

It supports compliance with applicable law and regulation — including the EU AI Act, UK GDPR, and, where relevant, the FCA's expectations and SM&CR requirements. It complements rather than supersedes the organisation's Data Protection Policy, Information Security Policy, and sector-specific obligations.

1.2 Definition of an AI system

TEMPLATE DEFINITION — ADAPT AS REQUIRED

An AI system is any machine-based system that processes inputs — including data, content, and environmental information — to generate outputs such as predictions, recommendations, decisions, or content, using techniques including machine learning, deep learning, natural language processing, or large language models. This includes AI capabilities supplied by third-party vendors and embedded within software-as-a-service applications.

1.3 Organisational scope

This policy applies to:

- All employees, contractors, and third parties acting on behalf of [Organisation Name]
- All legal entities within the [Organisation Name] group: [list entities]
- All business functions and departments
- Both internally developed AI systems and AI tools supplied by third parties, including AI embedded in SaaS platforms

1.4 Exclusions

DRAFTING GUIDANCE

Document any AI systems or use cases excluded from this scope, with reasons. Examples: AI operated by regulated subsidiaries under a separate group policy; AI used solely for R&D under a distinct research governance framework.

SECTION TWO

2 Risk classification.

All AI systems in scope must be classified against the following tiers before use. Classification determines the governance requirements that apply. Teams self-assess their use case against these tiers as part of the approval process.

	DESCRIPTION	EXAMPLES	GOVERNANCE REQUIREMENTS
TIER 1 High risk	Consequential decisions affecting individuals, regulated outputs, or safety-critical functions.	Credit scoring, HR screening, clinical decision support, automated enforcement.	Full approval process, ongoing monitoring, audit trail, mandatory human oversight, periodic performance review.
TIER 2 Medium risk	Operational use with limited direct impact on individuals or regulated outcomes.	Internal forecasting, document classification, customer segmentation, fraud detection.	Simplified approval, periodic review, incident reporting requirements.
TIER 3 Low risk	Productivity tools, internal assistants, and low-stakes automation.	AI writing assistants, meeting transcription, internal search, code assist.	Registration in the AI register, acceptable-use conditions, annual review.

ISO **EU AI ACT ALIGNMENT**

Tier 1 High Risk systems should be reviewed against the EU AI Act Annex III list of high-risk AI applications. Systems falling within Annex III that are placed on the market or put into service from **2 August 2026** are subject to additional mandatory requirements under Article 17 — including technical documentation, conformity assessment, and registration in the EU database.

3 SECTION THREE

Accountability & governance.

3.1 Policy ownership

Policy owner	[Role — e.g. CISO / Chief Risk Officer]
Board / executive accountability	[e.g. Board Risk Committee / AI Governance Committee]
Business unit accountability	[Defined per function — document in Annex A]
Technical accountability	[e.g. Head of Data / Head of Engineering]
Day-to-day administration	[e.g. GRC Team / Compliance Team]

3.2 Decision rights matrix

All approval authorities must be documented by name in the AI register at the point of approval.

DECISION	AUTHORITY	ESCALATION PATH
Approve Tier 1 (High Risk)	[AI Governance Committee / CRO]	[Board Risk Committee]
Approve Tier 2 (Medium Risk)	[CISO / Head of Compliance]	[CRO]
Approve Tier 3 (Low Risk)	[Business Unit Head / Line Manager]	[CISO]
Suspend AI system pending investigation	[CISO / CRO]	[CEO / Board]
Escalate AI governance concern	[Any employee — via channel]	[GRC Team / Compliance Hotline]
Decommission AI system	[As approval authority for tier]	[Document reason in AI register]

4

SECTION FOUR

AI use approval process.

Any AI system in scope must complete the following process before deployment. Requirements are differentiated by risk tier.

1

Initial triage

Classify the system against the Section 2 tier definitions. Complete the self-assessment form. Output: a completed triage record with tier assigned.

2

Impact assessment

Document intended use; affected individuals or decisions; data inputs and sources; potential failure modes and harms; whether human oversight is required.

Tier 3: short self-assessment · Tier 1: full AI impact assessment

3

Legal & compliance review

Confirm alignment with UK GDPR (including Article 22 where applicable), EU AI Act obligations, sector regulation (FCA, PRA, ICO), and employment law where AI affects HR decisions.

Tier 3: checklist sign-off · Tier 1: legal team sign-off

4

Security & data review

Confirm data handling and storage, access controls, vendor due diligence for third-party AI, and retention and deletion provisions. Link to a DPIA where a Tier 1 system processes personal data.

5

Approval decision

Formal approval by the authority defined in the Decision Rights Matrix (3.2). The decision must be documented with rationale, conditions of use, and the approver's name and date.

6

Conditions of use

Document any restrictions: prohibited use cases, mandatory human-review requirements, permitted and excluded data inputs, and user training requirements.

7

Registration

Add the approved system to the AI register (Section 5) with all required fields completed. Deployment may not proceed until registration is confirmed.

5

SECTION FIVE

AI register.

[Organisation Name] will maintain a central AI register as the single source of truth for all AI systems in use. Maintenance is the responsibility of [GRC Team / Compliance Team]. The register must be reviewed quarterly and whenever a system is approved, materially changed, or decommissioned.

Required register fields

FIELD	DESCRIPTION	MANDATORY?
System name & version	The AI tool or system name and version in use	Yes
Business unit / owner	The team and named individual accountable for this system	Yes
Risk tier	Tier 1 / 2 / 3 as assigned at approval	Yes
Approval date	Date formal approval was granted	Yes
Approved by	Name and role of the approving authority	Yes
Conditions of use	Summary of restrictions or mandatory oversight requirements	Yes
Scheduled review date	Date of the next governance review	Yes
Current status	Active / Under review / Suspended / Decommissioned	Yes
Linked DPIA	Reference to DPIA where personal data is processed	If applicable
Third-party vendor	Vendor name and contract reference for third-party AI	If applicable
Deregistration date	Date the system was retired, with reason	On retirement

SECTION SIX

6 Audit and monitoring.

6.1 What must be logged

TIER 1 & 2

- Decisions made by or with AI assistance — inputs, outputs, and confidence levels where available
- Human review actions on AI-generated recommendations
- Model updates, retraining events, and training-data changes
- Data changes that could affect model outputs
- User access events

ALL TIERS

- System deployment and decommissioning events
- Policy exceptions and out-of-policy uses
- Incident reports and near-misses

6.2 Retention

Audit logs must be retained for a minimum of [X years], consistent with the organisation's data retention schedule and applicable legal requirements. For regulated AI systems, framework-specific retention requirements take precedence.

6.3 Review frequency

RISK TIER	REVIEW FREQUENCY	REVIEWED BY
Tier 1	Quarterly	[AI Governance Committee]
Tier 2	Bi-annually	[CISO / Compliance Team]
Tier 3	Annually	[Business Unit Owner]

6.4 Out-of-cycle review triggers

- Significant model update or change to training data
- Evidence of data drift or performance degradation
- Complaints or disputes relating to AI outputs
- Regulatory guidance updates from the ICO, FCA, or EU bodies
- A material AI incident (see Section 7), or new deployment of a Tier 1 system

7 SECTION SEVEN

AI incident response.

7.1 What constitutes an AI incident

AI incidents include, but are not limited to: discriminatory or biased outputs; unexplained or contested decisions; personal data exposure caused by AI outputs; evidence of model manipulation or adversarial attack; significant performance degradation; sustained outputs outside expected parameters; and near-misses where any of the above were averted.

7.2 Incident reporting

Reporting channel	[GRC platform / compliance inbox / named contact]
Initial report within	[e.g. 24 hours of identification]
Reported to	[CISO / Compliance Team]
Regulatory notification	UK GDPR: 72 hours to the ICO · FCA / PRA: per applicable rules · EU AI Act: per Article 73

7.3 Response protocol

Contain	Assess whether the system should be suspended pending investigation. Preserve evidence — do not alter logs or outputs.
Investigate	Identify root cause. Document inputs, outputs, and any affected individuals. Assign a named investigation lead.
Notify	Notify affected parties where required under UK GDPR or applicable regulation. Notify regulators within required timeframes.
Remediate	Address root cause. Document changes made. Obtain sign-off before redeployment.
Review	Conduct a post-incident review. Document findings and lessons. Update the AI register and this policy if required.

8

SECTION EIGHT

Employee responsibilities & acceptable use.

8.1 Employee responsibilities

All employees using AI tools are responsible for:

- Using only AI tools approved through the process in Section 4
- Complying with the conditions of use attached to each approved system
- Reporting suspected AI incidents or policy breaches via the Section 7 channel
- Completing AI awareness training as required by [Organisation Name]
- Not using personal AI tools for work purposes without prior approval

8.2 Prohibited uses

The following are prohibited without prior written approval from [AI Governance Committee / CISO]:

- Entering personal data, confidential information, or client data into publicly accessible AI tools
- Using AI to make solely automated decisions with legal or similarly significant effects on individuals, without the human-review provisions required by UK GDPR Article 22
- Deploying AI in safety-critical functions without Tier 1 approval and mandatory human oversight
- Using AI to generate content that could be misleading, discriminatory, or in breach of applicable law
- [Add organisation-specific prohibited uses]

8.3 Training

Requirement	[e.g. Annual AI awareness module]
Owner	[e.g. HR / L&D / Compliance]

8.4 Consequences of breach

Breaches will be handled under [Organisation Name]'s disciplinary procedures. Material breaches may result in suspension of access to AI systems pending investigation.

9 SECTION NINE

Policy review.

9.1 Scheduled review

This policy will be reviewed at a minimum annually. The policy owner is responsible for initiating the review and coordinating input from legal, compliance, risk, and technology functions.

Next scheduled review	[Date]

9.2 Out-of-cycle review triggers

This policy must be reviewed out of cycle when:

- Significant new AI systems are deployed across the organisation
- Relevant legislation or regulatory guidance changes (EU AI Act, ICO, FCA, PRA)
- A material AI incident reveals gaps in the policy
- The organisation undergoes material structural change affecting AI governance
- The AI Governance Committee flags emerging use cases not covered by current scope

9.3 Review process

Reviews must be documented with the date of review; who conducted it; what changes were made and why; and confirmation of approval by the policy owner. Changes take effect from the version publication date. All staff with AI system responsibilities must be notified of material changes.

A FIXED ANNUAL REVIEW IS NOT ENOUGH ON ITS OWN

Given the pace of AI regulatory development, the out-of-cycle triggers above are as important as the scheduled review. Treat regulatory change as a live input, not an annual checkpoint.

BEFORE YOU CIRCULATE

First-draft review checklist.

Use this before circulating the policy for approval. These are the gaps most commonly identified in first-draft AI governance policies.

- [ ] **Scope covers third-party AI tools, not only internally developed systems**
  Most AI risk sits in vendor-supplied tools — policies scoped to internal AI only miss the majority of exposure.
- [ ] **An AI system is defined in the policy**
  Without a definition, teams cannot determine what is in scope.
- [ ] **An AI register is required and its fields are specified**
  Without a register, there is no single view of what AI systems are approved and in use.
- [ ] **Accountability roles have decision rights attached**
  Naming an oversight body is not the same as specifying who can approve, escalate, or suspend.
- [ ] **The approval process produces a documented output**
  An assessment form with no attached decision record does not constitute governance evidence.
- [ ] **The incident definition covers gradual failures and disputed outputs**
  AI failures often accumulate slowly — a definition borrowed from IT security will miss them.
- [ ] **Tier 1 systems are cross-referenced to a DPIA requirement**
  UK GDPR Article 35 requires a DPIA for high-risk processing — including many AI applications.
- [ ] **Review triggers include regulatory changes and material incidents**
  A fixed annual review is insufficient given the pace of AI regulatory development.
- [ ] **SM&CR mapping is completed for regulated financial activities**
  Named Senior Managers bear personal accountability where AI influences regulated activities.
- [ ] **A third-party AI / vendor risk management reference is included**
  Vendor-supplied AI requires governance — document how third-party AI is assessed and approved.

PUT IT INTO PRACTICE

A policy is the governance layer. The controls sit behind it.

An AI inventory in practice, pre-deployment assessments, and the audit-trail structures regulators expect are what turn this document into operational governance.

IMPLEMENTATION GUIDE	**How to Build an AI Governance Policy**	surecloud.com/resource-hub
COMPLIANCE FRAMEWORK	**ISO 42001:2023 — AI Management System**	surecloud.com/frameworks/iso-42001
REGULATORY LANDSCAPE	**AI Governance Regulations: UK & EU 2026**	surecloud.com/blog-hub
PLATFORM	**Compliance Management on SureCloud**	surecloud.com/product

See SureCloud in action.

Gracie AI Agents and Skills automate the evidence collection, monitoring, and audit-trail work your AI governance policy demands — at scale.

Book a personalised demo →