

IOTA

AI Governance Framework

Dr Georgiana Marsic
Head of Data and AI Assurance
Lloyds Banking Group, UK



Overview

- **Definitions**
- **Architecture of an AI Governance Framework**
- **Charter**
- **Accountability and Control Environment**
- **Lifecycle Governance**
- **Maturity and Adaptive Evolution**



Definitions

Core concepts

Trustworthy AI

The outcome: AI worthy of trust (people perceive it as reliable, safe, secure, fair, private, transparent and explainable).

AI Ethics

Moral principles and human values that guide AI development, deployment and use. The principles that determine whether AI is worthy of trust.

Responsible AI

The practice of developing and using AI in a way that benefits society while minimising the risk of negative consequences. Covers translating principles into deployable practices.

AI Governance

Structures, policies, processes, roles, decision rights, accountability and oversight put in place to oversee and manage AI in an organisation.

AI Compliance

Demonstrable conformity with binding requirements (laws, regulations, contracts, standards) and organisation policies.

Architecture of an AI Governance Framework

The CALM AI Governance Framework

C

Charter

AI Governance Purpose · Ethics principles · Scope and boundaries · Intended outcomes

A

Accountability & Control Environment

Roles & decision rights · Governance structures · Risk management · Policies & processes · Controls · Assurance

L

Lifecycle Governance

Intake & inventory · Impact & risk assessment · Pre- and post-deployment governance · Change control · Incident management

M

Maturity: Capability & Adaptive Evolution

Skills & literacy · Adapting to change · Platforms, tools & guardrails · Operations · Continuous improvement

Charter

Core components

The Charter is the foundation of AI governance: it sets why it exists, the principles that guide it, what it applies to, and the direction and outcomes it enables.



AI Governance Purpose

Why do we need AI governance?

Why the organisation needs AI governance: its role in enabling trustworthy, responsible and value-creating AI adoption.



Ethics Principles

What principles guide our decisions?

The values that guide responsible AI decisions:

- Accountability · Fairness · Reliability · Transparency
- Privacy · Safety · Human oversight · Contestability



Scope & Boundaries

What is in scope, and where are the limits?

What governance applies to: AI systems, use cases, third-party & GenAI / agentic AI, data, users, business areas and lifecycle stages.

- Boundaries: in scope · out of scope · restricted · prohibited



Strategic Direction & Intended Outcomes

Where are we heading, and what should it enable?

Leadership intent and priorities for AI adoption, and the outcomes governance is designed to enable:

- Responsible innovation · risk-managed adoption
- Regulatory readiness · stakeholder trust
- Sustainable value · better customer / employee / citizen outcomes

CHARTER

AI Governance Purpose

Why do we need AI governance? Different motivations, one purpose

Governance exists so the organisation can adopt AI it can trust, defend and scale.

CONSEQUENCE

+55%

rise in documented AI incidents between 2024 and 2025

THE CASE

AI failures harm people and create real liability.

WHAT GOVERNANCE DOES

Prevent harm: know what AI runs, who owns it, when to stop it.

TRUST

83%

would be more willing to trust AI when assurance mechanisms are in place

THE CASE

Trust is the condition for adoption and value.

WHAT GOVERNANCE DOES

Earn trust: show evidence that AI is reliable, fair and accountable.

REGULATORY

€35M / 7%

maximum EU AI Act fine for prohibited AI practices (Art. 99)

THE CASE

Regulation is accelerating. Compliance is the floor.

WHAT GOVERNANCE DOES

Meet obligations: classify risk, apply controls, keep evidence.

STRATEGIC

72%

of S&P 500 companies flagged AI as a material risk in 2025, up from 12% in 2023

THE CASE

Investors price AI risk. Governed AI scales faster.

WHAT GOVERNANCE DOES

Scale with confidence: clear approval paths and risk limits.

AI GOVERNANCE PURPOSE

Enable trustworthy, responsible and value-creating AI adoption: prevent harm, earn trust, meet obligations and scale with confidence.

Charter: AI ethics principles

The values that shape real AI choices

Accountability

Named individuals bear responsibility for AI outcomes. Audit mechanisms, assessments and redress pathways must be in place before deployment.

Fairness & Non-discrimination

AI must not discriminate by design or by effect. Training data, design choices and deployment contexts must all be audited for bias and exclusion.

Professional Responsibility

Those who build and deploy AI bear a duty of care to design, build and use AI responsibly, and anticipate long-term consequences.

Privacy

AI must collect and use only what is needed, use data with consent, and protect individuals' data rights throughout the AI lifecycle.

Transparency & Explainability

People must know when they interact with AI or AI influences decisions. Decisions and logic must be understandable to those they affect.

Promotion of Human Values

AI should make human lives better in ways that are equitable, dignified and consistent with what we collectively care about.

Safety & Security

Systems must perform reliably as intended, resist external threats, and be tested against failure scenarios continuously, not just at launch.

Human Control

Humans retain the right to review, override and opt out of automated decisions. AI supports human judgment, it does not replace it.

Reliability & Robustness

AI systems must perform as intended. They should be built and assessed for accuracy and proper functioning across a wide range of conditions.

From principles to practice

Principles only create impact when translated into decision-guiding commitments, and then acted on (e.g. a named AI Owner, a path to remedy).

Charter: scope

Define the scope of AI Governance

What AI governance applies to, and where the organisation draws the boundary

SCOPE AREAS

	Builds or develops	Internal models, custom AI tools
	Buys or licences	Vendor AI, API-accessed foundation models
	Embeds in enterprise tools	AI capabilities within platforms, e.g. Microsoft Copilot, Salesforce Einstein, SAP BTP
	Permits employees to use	Approved tools, experimentation, informal Gen AI, e.g. ChatGPT, Claude
	Uses to make decisions	Hiring, credit, benefits, clinical support
	Deploys via third parties	Outsourced services, partner-operated AI-enabled processes & supply-chain AI

FOUR SCOPE DIMENSIONS

- 1 Organisational**
Which entities, business units and third parties are subject to the framework
- 2 Systems and assets**
Which systems, assets and AI models are within scope
- 3 Process**
Which decision types, lifecycle stages and change events trigger governance
- 4 People**
Which roles, functions and external parties, including vendors and regulators, are bound

Charter: strategic direction & intended outcomes

The trajectory & what governance should enable

Leadership intent for AI governance scaled to ambition and consequence.



Increasing ambition, trust and value → governance grows with the trajectory you choose

The outcomes AI governance is intended to enable:

Responsible innovation

Risk-managed adoption

Regulatory readiness

Stakeholder
trust

Scalable operating model

Better customer / employee / citizen
outcomes

Sustainable value

Accountability & Control Environment

Two core components

The control environment defines who is accountable, how AI decisions are governed, and how AI risks are managed, controlled and assured.



ACCOUNTABILITY & CONTROL ENVIRONMENT

Creates the organisational conditions for responsible, controlled and trustworthy AI adoption.



Accountability & Decision Governance

Who owns, decides, approves and escalates AI-related decisions.

- Senior accountability for AI governance and outcomes
- AI system / use case owners
- Decision rights and delegated authority
- Governance forums and approval bodies
- Risk-area owners: Legal, Privacy, Security, Ethics, Architecture
- Operating model, escalation routes and cross-functional challenge



Risk, Control & Assurance System

How AI risks are identified, managed, controlled, monitored and evidenced.

- AI risk taxonomy and risk appetite
- Use case intake, triage and classification
- Impact and risk assessment process
- Policies, standards, procedures and playbooks
- Preventive, detective and corrective lifecycle controls
- Control evidence, monitoring, reporting and assurance reviews

Accountability & decision governance

Decision rights across the AI lifecycle

Decision rights define who can approve, reject, escalate, pause, change or retire an AI system at each stage.

1 Set the rules

AI policy, risk appetite, prohibited & acceptable use, standards

2 Classify the use case

AI definition, internal risk tier, regulatory category, affected stakeholders

3 Approve data, model & supplier

Data permissions, privacy, security, vendor risk, model suitability

4 Accept the risk

Impact assessment, mitigations, residual risk, exceptions

5 Approve deployment & change

Validation, go-live, monitoring thresholds, material-change approval

6 Intervene when needed

Escalation, incident response, rollback, suspension, retirement

Accountability architecture

Board oversight

Sets risk appetite; demands assurance; holds executives to account.

Executive accountability

A named executive answers for the use case.

Risk / control challenge

Second-line functions challenge the evidence.

Technical / system owner

Owens how the system is built and operated.

Business / product owner

Owens the use case and its outcomes.

Plus an independent assurance route across all layers.

Key examples of AI Governance body types

From decision rights to governance bodies

Match structure to decision volume and risk.

Body	Primary role	Typical authority
AI Governance Council	Enterprise oversight, risk appetite, policy ratification, high-risk escalation	Approve, reject, set appetite, escalate to executive / board
AI Steering Committee	Operational coordination, prioritisation, deployment decisions	Approve medium-risk use, resolve conflicts, escalate high-risk
AI (& Data) Ethics / Review Board	Ethical challenge, fairness, human impact, novel-risk review	Advisory, escalation or veto depending on mandate
Smaller orgs may merge Council and Steering Committee, but keep an accountable owner, an escalation path, and a documented decision record.		

AI governance is a team sport: accountability spans the whole organisation

Board	Senior leadership	Business executives	Responsible AI team
Risk Management	Legal	Compliance	InfoSec / Cyber
Digital Architecture	Data Governance	Data Privacy	Data Science & AI
CIO, CTO, CDO / CDAIO	Business Ethics	HR	Procurement
Product Leaders	Marketing / Comms	Gov't relations	Internal Audit

Risk, control and assurance system

Governance as competitive advantage

If any link in this chain is missing, the control environment breaks.

Risks

Identify what could go wrong, who could be harmed, and the level of risk the organisation will accept.

Policies

Convert risk appetite and principles into mandatory obligations and boundaries.

Processes

Define the repeatable routes through which those obligations are executed.

Controls

Embed specific checks, gates and actions into processes to prevent, detect or correct failure.

Evidence

Creates the record that decisions were made and controls operated as required.

Assurance

Tests whether the controls are well-designed, operating effectively and trusted.

Risk management: the risk universe

AI risk is not a new universe

AI risks amplify and connect the risk domains the enterprise already manages.

RISK MANAGEMENT

The systematic identification, assessment and treatment of uncertainties that could affect objectives. It sets risk appetite, selects treatments, and monitors residual risk.

TREATMENTS

Avoid

Reduce

Transfer

Accept

AI governance extends, not replaces, enterprise risk management.



AI risks

Adapted from *Governing the Machine: How to Navigate the Risks of AI and Unlock Its True Potential* by Ray Eitel-Porter, Paul Dongha & Miriam Vogel.

1 Accuracy & Reliability

Correct, stable, dependable outputs

2 Fairness & Bias

No disparate impact on protected groups

3 Interpretability & Transparency

Explainable, understandable decisions

4 Accountability

Clear ownership of outcomes

5 Privacy

Lawful, minimised use of personal data

6 Security

Resilient to misuse and attack

7 IP & Confidentiality

Protect proprietary data and rights

8 Workforce


Impact on people, skills and roles

9 Environmental & Sustainability

Compute, energy and footprint

Proportionality

Example: EU AI Act risk tiering



Risk tiering

Risk tier is according to the outcome

Focus on the potential impact on people

THE LABEL SAYS

“Decision support”

Sounds low-stakes.
Sounds advisory.



BUT IT ACTUALLY

- ▶ Ranks
- ▶ Filters
- ▶ Rejects
- ▶ Influences high-stakes outcomes

Look at what the system does to real people, not the system description / label.

An AI system that "supports" recruiters by auto-rejecting candidates is making high-stakes decisions that impact people's lives.

Internal risk tier vs regulatory category

Internal risk tier

Drives control depth, approval route, escalation, monitoring and assurance intensity.

Regulatory category (EU AI Act)

Drives legal obligations, conformity assessment, compliance evidence and reporting duties.

Importance of considering both lenses

A customer-facing GenAI assistant may be only Limited-risk under the Act (“just a chatbot”), yet High internal tier because it reads confidential client data and writes customer communications.

AI Governance processes

The different AI Governance processes that may exist in organisations

Grouped by the NIST AI RMF functions: Govern wraps Map, Measure and Manage across the lifecycle.

GOVERN

Governance structures & policy

Data governance

Vendor / TPRM

MAP

find & frame

AI Intake, Triage & Inventory

Risk Assessment

Impact Assessment

MEASURE

test & evaluate

Model Validation & Testing

Bias & Fairness Testing

Safety, Security & Red-teaming

MANAGE

control, operate & respond

Controls & Documentation

Monitoring & Reassessment

Agentic AI Authorisation & Runtime Oversight

AI intake, triage and inventory

The foundation every other process depends on

How AI enters governance, gets screened and routed, and becomes a living record

1 · INTAKE

entry point

- Manual submission
- New build, buy or SaaS request
- Automated discovery → shadow AI
- SaaS / expense / SSO / code scans
- Procurement & vendor trigger
- Change, attestation & incident

2 · TRIAGE

screen & route

- In scope? Based on AI-definition
- New or already registered?
- EU AI Act role (Provider/Deployer/etc)
- Provisional risk screen + route trigger
- Assign accountable owner & record
- Route & trigger assessments → fast-track · full review · prohibit

3 · INVENTORY

the living register

- Purpose & business owner
- Build / buy / SaaS + vendor
- Data sources & sensitivity
- Risk tier & EU AI Act category
- Approval & monitoring status
- Lifecycle status · maintained continuously

Inventory is the persistent foundation: assessment, approval and monitoring read from and write to it, and every material change re-enters intake.

AI risk & impact assessment

Two lenses: risk TO the business, and impact ON people

Some run them separately; others combine them.

AI RISK ASSESSMENT

Focus: risk TO the organisation

- › Internal risk tiering (likelihood × consequence)
- › Model, data, security, operational, vendor & legal risk
- › Control requirements & residual risk vs appetite
- › Regulatory exposure & approval routing

Typically owned by: Risk / second line

AI IMPACT ASSESSMENT

Focus: impact ON people & society

- › Affected individuals & groups; vulnerable users
- › Fairness, bias & discrimination
- › Fundamental rights (EU AI Act FRIA), transparency, contestability
- › Societal & reputational consequences

Typically owned by: Ethics / Privacy / Legal

COMBINED

One “AI Risk & Impact Assessment” with a single intake routing to more detailed assessments: lower friction.

Control

Process → Control Mapping

A control is a specific activity, owned by someone, embedded in a process and producing evidence.

Process	Preventive	Detective	Corrective
Intake, Triage & Inventory	Mandatory intake before procurement	Inventory reconciliation	Escalate unregistered use
AI Risk & Impact Assessment	Assessment before approval	Review incomplete assessments	Rework or reject
Vendor / Third-party Risk Management	Due diligence before contract	Vendor-change monitoring	Suspend / remediate
Monitoring	Monitoring plan before go-live	Output / usage monitoring	Reassessment / suspension

Three lines of assurance

first line owns & operates · second line reviews, challenges & verifies · third line independently tests design & operating effectiveness.

Assurance

Activities that turn controls into evidence-based confidence

ASSURANCE

The structured set of activities that establish, evaluate and provide **evidence-based confidence** that controls are appropriately designed, implemented, operating effectively and managing risks within appetite.

NOT THE SAME AS OPERATING CONTROLS


Assurance evaluates whether the control environment works as intended. It does not replace the activities that operate the controls, or general risk management.

Interconnected, not necessarily sequential.



The AI lifecycle is a governed loop

Every stage in the loop has a named owner



Accountability

Business owner	Accountable for the AI solution itself: purpose, use case, outcomes.
Technical / model owner	Operationally accountable: performance, monitoring execution, model behaviour.
Risk / legal / privacy / security	Challenge, advise, gate. Accountable for the review, not for the solution.
Approval authority / committee	Decides based on risk tier. The formal governance decision point.
Audit	Tests whether process and evidence exist. Independent of the approval chain.

Governance happens at every stage

Not just before launch

Each stage has a gate question. The answer must be evidenced before the system progresses.

Lifecycle stage	Gate question	Evidence / decision output
Intake	Is this AI, and does it need governance?	Intake record + inventory entry
Assessment	What risk tier and specialised reviews / approvals apply?	Impact & risk assessment
Design	What controls must be built in?	Control plan + oversight design
Development	Has it been tested against intended use?	Validation, bias/fairness, security evidence
Deployment	Is approval evidence complete?	Signed approval + release decision
Monitoring	Is it behaving as expected in production?	Metrics, thresholds, owner review
Change	Has anything material changed?	Reassessment / reapproval decision
Incident	Has harm, misuse or failure occurred?	Containment + incident record
Retirement	Can it be safely decommissioned?	Archived evidence + closed obligations

The AI inventory

You cannot govern what you cannot see

Unique ID + AI solution name

Use case description (what it does)

Business owner (named)

Technical owner

Data sources & inputs

Vendor / third-party flag

Risk tier

Regulatory classification

Lifecycle stage

Monitoring status

Last governance review

Reapproval trigger date

Autonomous action Yes / No, and if Yes, the documented action scope.

Risk and impact assessment

Two lenses: risk TO the business, and impact ON people

1 · Impact on people

Who is affected (employees, customers, public, vulnerable groups)? Rights, opportunities, access? Bias / discrimination risk? Does it override human judgment?

2 · Regulatory risk

EU AI Act class; GDPR Art. 22 automated decisions & special-category data; sector rules; cross-border / jurisdiction flags.

3 · Operational risk

Cost of error & reliability; explainability to individuals and regulators; dependency; data-quality / drift exposure.

4 · Reputational / strategic

Public sensitivity; media & civil-society scrutiny; alignment with stated organisational values.

Agentic overlay · autonomy · reversibility · traceability

RISK TIER

The tier is the routing decision: it sets the entire governance track.

Three tiers: three governance tracks

Tier	Profile	Governance intensity
Tier 1 · HIGH	High-risk (EU AI Act) · material people impact · autonomous action with irreversible consequences · customer-facing · material operational impact	Full assessment · mandatory approval gate · independent review · continuous monitoring · enhanced incident response
Tier 2 · MEDIUM	Moderate people impact · sector regulatory triggers · supervised or limited autonomy · use of personal data · material IP impact	Specialised reviews · approval gate · periodic monitoring · standard incident response
Tier 3 · LOW	Minimal people impact · no regulatory triggers · no autonomous action · no IP impact	Lightweight assessment · inventory registration

Agentic risk depends on what the system can do

- 1 Can it act without human approval?
- 2 Can it access tools, money, records or comms channels?
- 3 Can its actions be reversed?
- 4 Can we explain why it acted, after the fact?
- 5 Can we stop it immediately?

DEFAULTS TO TIER 1 HIGH RISK

Material autonomy, irreversible action, or limited traceability, regardless of action type.

Routing the use case to specialist review

One use case can trigger several reviews in parallel

Specialist review	Routed when...	Owner → output
Privacy / data protection	Personal or special-category data; large-scale monitoring; automated decisions (GDPR Art. 22)	DPO → DPIA
Fundamental-rights impact	EU AI Act high-risk use; public services; impact on vulnerable groups	AI Governance + Legal → FRIA
Legal & regulatory	Regulated or high-risk use; new jurisdiction; liability or enforcement exposure	Legal → regulatory assessment
Intellectual property	Training on licensed/copyright data; generative IP impact; trade-secret exposure	Legal (IP) → licensing opinion
Security & cyber	Tool / API / system access; credentials; external connectivity; autonomous actions	InfoSec → threat & access review
Ethics & fairness	Discrimination or bias risk; sensitive or public-facing use; value-laden decisions	Ethics panel → fairness review
Procurement / third-party	Vendor, embedded or third-party model; external data sharing; vendor model-change	Procurement (TPRM) → vendor AI assessment
Model risk & validation	Material decision models: credit, pricing, underwriting; financial-services models	Model risk → independent validation
HR / employment	Workforce screening, ranking, monitoring or performance use	HR + legal → employment review

Maturity: Capability and Adaptive Evolution



Governance is one of four HOW pillars representing the strategy delivery system, the system that turns AI strategy into trusted execution. It cannot work alone. It depends on the other three..

What Governance Needs to Work

Capability: Governance without these three pillars is just a set of rules nobody follows.

People & Culture

What it provides:

The right people with the right skills, and a culture where doing the right thing is normal.

Without it:

Rules exist on paper but nobody has the time, training, or confidence to follow them.

Examples:

Named governance roles, AI literacy training, psychological safety to speak up

Operations

What it provides:

The day-to-day processes, performance measures, and ways of working that embed governance into routine.

Without it:

Governance only appears at project kick-off and after disasters. It's a fire alarm, not a heartbeat.

Examples:

Gate reviews, monthly dashboards, incident response drills, improvement cycles

Digital Capabilities

What it provides:

The platforms, tools, data architecture, and technical standards that make governance enforceable.

Without it:

Decisions are untracked, evidence is lost, and the organisation cannot prove compliance.

Examples:

AI system inventory, audit logs, monitoring dashboards, data classification tools

Pillar 1: People & Culture

The Right People & Skills

- ✓ **Named roles with real authority**
Not just responsibility added to an already full job.
- ✓ **Skills mapped to gaps**
What can we do vs. what do we need?
Reviewed yearly.
- ✓ **Backup people for key roles**
One person leaving should not collapse governance.
- ✓ **Dedicated time allocated**
Governance cannot be a side-of-desk activity.

The Right Culture

- ✓ **Leaders model the behaviour**
If the board ignores governance, everyone else will too.
- ✓ **Safe to speak up**
People must feel they can say 'STOP' without punishment.
- ✓ **AI literacy at every level**
Board: oversight questions. Staff: AI use and guardrails
Everyone: basics.
- ✓ **Governance celebrated, not resented**
When a gate stops a bad deployment, that is a success.

Pillar 2: Operations

Operating Model

How governance runs day-to-day. Not a document on a shelf, but the rhythm of how governance teams engage with business and technology at defined touchpoints, with authority to intervene.

- ✓ **Governance checkpoints**
Built into the development lifecycle, not bolted on afterwards.
- ✓ **Embedded governance partners**
Governance people sitting inside business units, not in an ivory tower.
- ✓ **Monthly operating rhythm**
Risk register review, controls check, escalation triage. Every month.
- ✓ **Governance as a service**
For teams without their own specialists: accessible help, not distant bureaucracy.

Performance Management

How we know governance is actually working. The critical distinction: measure outcomes (risks prevented, harms avoided), not activity (meetings held, forms completed).

- ✓ **Gate compliance rate**
% of changes through correct gates. Target: 100%. Shortfall = root cause review.
- ✓ **Time-to-decision**
How long from gate submission to decision. Too slow = pressure to bypass.
- ✓ **Incident response speed**
% of incidents triggering the right protocol within required timeframe.
- ✓ **Training completion**
% of governance role-holders with current training. Below 90% is a red flag.

Pillar 3: Digital Capabilities

Governance Tooling & Platforms

Registers, workflows, evidence repositories, and dashboards that make governance consistent, auditable, and scalable.

- ✓ AI system register and data catalogue
- ✓ Policy and obligation register
- ✓ Gate workflow management (not email!)
- ✓ Immutable audit logs for all decisions

Data & Information Architecture

How data flows, who owns it, and how it supports governance decisions. AI governance is only as good as the data it relies on.

- ✓ Data classification (public / confidential / restricted)
- ✓ Data minimisation enforced by architecture
- ✓ Data lineage: source to model to output
- ✓ Named data stewards for each domain


Solution Architecture & Infrastructure

The technical standards all systems must meet, and the infrastructure they run on. Architecture decisions ARE governance decisions.

- ✓ Human override built into every AI system
- ✓ Model versioning and tested rollback
- ✓ Explainability requirements for high-risk AI
- ✓ Patch management, DR testing, availability SLAs

Without these capabilities, governance becomes unscalable, evidence is lost, and compliance cannot be demonstrated.

Adaptive evolution: keeping the system alive



Continuous Improvement

What have we learned from operating the system that should change how any part of it works?



Horizon Scanning

What is coming next in technology, regulation, and risk that will require the strategy, governance, or capabilities to change?



Continuous improvement and horizon scanning are not specific to governance. They are the adaptive mechanism for the entire strategy blueprint. When the world changes, every pillar and every layer of the WHY-WHAT-HOW must respond.

IOTA

Thank You

