AI GOVERNANCE

STARTER KIT

Practical resources to establish accountability, assess AI risks, and support responsible AI adoption.

1. Introduction: Why AI Governance Matters

What is AI governance?

AI governance is the system of policies, roles, processes, and controls an organization uses to ensure that artificial intelligence is developed, procured, and used safely, ethically, and in line with legal and business requirements. It sits alongside — and borrows heavily from — existing disciplines like data governance, IT risk management, and information security, but adds AI-specific concerns: model behavior, bias, explainability, and the pace at which AI capabilities and regulations are evolving.

In practice, AI governance answers three questions for every AI system in use: Who is accountable for it? What risks does it introduce? And how do we know it's operating as intended over time?

Why organizations need AI governance

AI adoption is usually driven bottom-up — individual teams adopt chatbots, copilots, and automation tools long before any formal review happens. Without governance, organizations end up with AI systems making or influencing decisions that no one has formally reviewed, approved, or is accountable for.

- Regulatory exposure: laws like the EU AI Act and sector-specific rules increasingly require documented risk assessments, human oversight, and transparency for certain AI use cases.
- Reputational risk: biased hiring tools, inaccurate customer-facing chatbots, or opaque automated decisions can cause public harm quickly and publicly.
- Operational risk: AI systems can fail silently — a model can degrade in accuracy over time without anyone noticing unless it is monitored.
- Competitive advantage: organizations with mature governance can adopt new AI capabilities faster and with more confidence, because the guardrails are already in place.

Example: *A bank rolls out an AI-based loan pre-screening tool in one regional branch without a formal risk review. Six months later, regulators ask for documentation showing how bias was tested and how decisions can be explained to applicants — and none exists. Governance, applied from day one, would have required that documentation as a condition of deployment, not as an afterthought.*

Key challenges of unmanaged AI

- Shadow AI: employees using public AI tools with sensitive company or customer data, with no visibility from IT or security.
- Fragmented ownership: no single person or team can say how many AI systems are in use, who owns them, or what data they touch.

- Inconsistent risk treatment: a high-risk system (e.g., one influencing hiring or credit decisions) may receive the same light-touch review as a low-risk internal tool (e.g., a meeting-notes summarizer).
- Vendor blind spots: third-party AI features are embedded inside existing SaaS tools, often enabled by default, without any assessment of how they handle data.

How governance supports responsible AI adoption

Good governance is not a brake on innovation — it is what allows an organization to say yes to AI faster and with more confidence, because every new use case is evaluated against a known, repeatable process rather than starting from scratch. This starter kit gives you the core building blocks: an inventory, an accountability model, risk assessment tools, policy language, and a rollout plan you can execute in 30 days.

2. AI Governance Quick-Start Guide

These five steps form the backbone of every mature AI governance program. They are sequential — each step depends on the output of the one before it — but they are also cyclical: once you complete step 5, you return to step 1 as new AI systems are adopted.

Step 1 — Identify your AI systems

You cannot govern what you cannot see. Start by cataloguing every AI system in use across the organization, including tools embedded inside existing software (e.g., an AI writing assistant built into your CRM) and AI features employees have turned on themselves.

- Survey department leads and IT/procurement records for known AI tools.
- Check SaaS platforms for AI features that may already be enabled by default.
- Review browser and network logs for traffic to public AI services, where feasible and appropriate.

Step 2 — Assign ownership and accountability

Every AI system needs a named business owner (accountable for its use and outcomes) and a technical owner (accountable for how it's built, configured, and maintained). Without named owners, risk issues have nowhere to land.

Step 3 — Assess AI-related risks

Use a consistent method (Section 5 and 6 of this kit) to evaluate each system across privacy, security, bias, transparency, accuracy, compliance, human oversight, and third-party risk.

Step 4 — Establish policies and controls

Translate your risk findings into concrete rules: acceptable use policies, data handling requirements, approval workflows for new AI tools, and technical controls such as access restrictions or output monitoring.

Step 5 — Monitor and continuously improve

AI systems, vendor terms, and regulations all change. Build a cadence — quarterly is common — for re-reviewing your inventory, re-scoring risks, and updating policies.

Example: *A mid-size insurer runs this five-step cycle quarterly. In Q1 they discover 14 previously unknown AI tools during the inventory step; by Q3, that number drops to 2 new tools per quarter because business units now know to register new AI use cases as part of procurement — governance has become routine rather than a special project.*

3. AI System Inventory Template

The inventory is the single source of truth for every AI system in the organization. It should be a living worksheet — reviewed and updated on a regular cycle — not a one-time exercise. Below is the field structure, followed by a worked example.

Fields to capture

- AI system/application — the name of the tool or model.
- Business purpose — what problem it solves and for whom.
- Department/owner — who is accountable for its use.
- Data used — types of data the system accesses or processes.
- AI model/provider — the underlying model or vendor (e.g., internal model, or a named vendor).
- Users — who interacts with the system, internally or externally.
- Risk level — Low / Medium / High / Critical, from your risk scoring exercise.

- Regulatory requirements — any laws or standards that apply (e.g., GDPR, EU AI Act, sector rules).
- Current controls — safeguards already in place.
- Review status — Not Reviewed / In Review / Approved / Needs Remediation.

Worked example

AI System	Purpose	Owner	Data Used	Model/Provider	Users	Risk	Reg. Reqs.	Controls	Status
Resume Screening Assistant	Ranks job applications	HR / Talent Acq.	Applicant CVs, PII	Vendor: TalentAI (SaaS)	HR recruiters	High	EU AI Act, GDPR	Human review of shortlists	In Review
Meeting Notes Summarizer	Summarizes meetings	Operations	Meeting audio/text	Internal licensed LLM	All staff	Low	Internal data policy	No external sharing; opt-in	Approved

Note how the two examples above receive very different treatment: the resume screener is high-risk because it materially affects people's employment outcomes and is explicitly in scope of AI-specific regulation, while the meeting summarizer is low-risk and can move through governance quickly.

4. AI Roles & Accountability Matrix

A RACI matrix clarifies who is Responsible (does the work), Accountable (owns the outcome), Consulted (provides input), and Informed (kept in the loop) for each governance activity. Ambiguity here is one of the most common reasons AI governance programs stall — everyone assumes someone else owns the risk.

Activity / Role	AI Inventory	Risk Assessment	Policy Approval	Incident Response
Executive leadership	I	I	**A**	**A**
AI governance team	**A**	**A**	R	R
IT / Security	R	R	C	R
Legal & compliance	C	R	R	C
Data teams	R	C	C	C
HR	C	C	C	I
Business owners	R	C	I	I
AI users	I	I	I	R

R = Responsible A = Accountable C = Consulted I = Informed

Example: *When a customer complains that a chatbot gave incorrect refund information, the RACI matrix tells you immediately who acts first (IT/Security and AI users file and triage the incident), who owns the resolution (the AI governance team), and who must ultimately answer for it externally (executive leadership) — instead of a scramble to figure out who's in charge.*

5. AI Risk Assessment Checklist

Run every AI system in your inventory through this checklist. The goal isn't a pass/fail — it's a structured set of answers that feed directly into the risk scoring template in Section 6.

Risk Area	Key Questions
Privacy	Does the system process personal or sensitive data? Is that data minimized, and do individuals have visibility or consent rights over its use?
Security	Can the system be exploited or manipulated — through prompt injection, data poisoning, model theft, or unauthorized access to outputs?
Bias	Could outputs create unfair outcomes for particular groups? Has the system been tested against representative data?

Transparency	Can decisions or outputs be explained to the people affected by them, in plain language?
Accuracy	How reliable are the AI outputs, and what happens when the system is wrong?
Compliance	Does the use case meet applicable legal, regulatory, and internal policy requirements?
Human Oversight	Can humans intervene, override, or stop the system when necessary — and is that pathway actually used?
Third Party	Are external AI vendors properly assessed for data handling, security, and reliability before adoption?

Example: *Applying this checklist to a customer support chatbot might reveal: it processes customer PII (Privacy — needs review), has no rate-limiting against prompt injection (Security — gap identified), was trained mostly on English-language transcripts (Bias — needs testing on non-English interactions), and has no human escalation path for angry customers (Human Oversight — gap identified). Each of these becomes a line item in the risk scoring and mitigation tracker.*

6. AI Risk Scoring Template

Once risks are identified, score them so you can prioritize. The simplest and most widely used model multiplies Likelihood by Impact to produce a Risk Score, which then maps to a risk category.

Scoring formula

Likelihood (1–5) × Impact (1–5) = Risk Score (1–25)

Risk Score	Category	Typical Response
1–4	Low Risk	Monitor periodically; document and move on.
5–9	Medium Risk	Assign an owner and mitigation timeline; review quarterly.
10–15	High Risk	Requires formal mitigation plan and leadership visibility before continued use.
16–25	Critical Risk	Immediate action; consider suspending use until mitigated.

Mitigation tracker fields

- Risk description
- Risk owner
- Mitigation action
- Target date
- Status (Open / In Progress / Mitigated / Accepted)

Example: *The resume-screening tool from Section 3 scores Likelihood = 4 (bias issues in hiring tools are common and well-documented) × Impact = 5 (affects people's livelihoods and carries regulatory exposure) = 20 → Critical Risk. This triggers an immediate mitigation plan: independent bias testing, mandatory human review of every shortlist, and a target date within 30 days — with the CHRO informed as risk owner.*

7. Responsible AI Checklist

A one-page reference teams can use before shipping or approving any AI-enabled feature or workflow.

☐ Fairness — outputs have been tested across relevant demographic and use-case groups for disparate impact.

☐ Transparency — users are told when they are interacting with AI, and how it influences decisions that affect them.

☐ Accountability — a named owner exists for the system's behavior and outcomes.

☐ Privacy — data collection and processing are minimized and disclosed to affected individuals.

☐ Security — the system has been assessed for adversarial misuse, unauthorized access, and data leakage.

☐ Human oversight — a person can review, override, or halt the system's outputs.

☐ Explainability — outputs can be described in terms a non-technical stakeholder can understand.

☐ Reliability — the system has been tested for consistent performance under realistic conditions.

☐ Continuous monitoring — a process exists to track performance and flag drift or degradation over time.

8. AI Governance Policy Starter Template

Below is a suggested structure for an organizational AI use policy, with example language you can adapt. This is a starting point, not legal advice — have counsel review the final version.

Purpose

"This policy establishes the principles, roles, and controls that govern the responsible development, procurement, and use of artificial intelligence at [Organization]."

Scope

Applies to all employees, contractors, and third parties who develop, deploy, or use AI systems on behalf of the organization, including AI features embedded in third-party software.

Acceptable AI use

- Use of approved AI tools for drafting, summarization, research, and internal productivity, subject to data handling rules below.
- Use of AI systems that have completed the risk assessment process appropriate to their risk level.

Prohibited AI use

- Entering confidential, customer, or regulated data into public/unapproved AI tools.
- Using AI outputs as the sole basis for decisions materially affecting individuals (employment, credit, eligibility) without human review.
- Deploying new AI-enabled features into production without a completed risk assessment.

Data protection

AI systems must comply with existing data classification and retention policies; personal data used in AI systems must have a documented lawful basis for processing.

Human oversight

High- and critical-risk AI systems must include a documented mechanism for human review, override, and escalation.

AI vendor management

New AI vendors must complete the third-party assessment in Section 11 before contracting, with results retained for audit.

Security requirements

AI systems must undergo a security review appropriate to their risk level, covering access control, data protection, and resilience to adversarial manipulation.

Monitoring and auditing

The AI governance team will re-assess the inventory and risk register on a defined cadence (recommended: quarterly) and report status to leadership.

Incident management

AI-related incidents (security, bias, accuracy, or compliance) follow the existing incident response process, with the AI governance team informed and involved.

Employee responsibilities

All employees using AI tools are responsible for following this policy, reporting suspected misuse or AI-related incidents, and completing required training.

9. AI Compliance & Framework Mapping

Rather than building governance activities separately for each regulation or standard, map your existing activities to the frameworks they already satisfy. This avoids duplicated effort and makes audits far easier.

Framework	How This Kit's Activities Map
NIST AI RMF	The Govern / Map / Measure / Manage functions correspond to Sections 4 (roles), 3 (inventory/mapping), 5–6 (risk assessment/scoring), and 12 (audit) of this kit.
ISO/IEC 42001	The AI management system requirements align with the policy template (Section 8), roles matrix (Section 4), and continuous improvement cycle (Sections 2 and 14).
EU AI Act	Risk-tiering (Section 6) mirrors the Act's risk categories; the inventory (Section 3) and human oversight checklist (Section 7) support required documentation for high-risk systems.
ISO/IEC 27001	Security requirements in the policy template (Section 8) and the security checklist (Section 10) extend existing information security controls to AI-specific risks.
GDPR	Privacy questions in the risk checklist (Section 5) and data protection policy language (Section 8) support lawful-basis and data-minimization obligations.
Organizational policies	All sections should be reconciled against existing internal policies (data governance, IT security, procurement) to avoid conflicting requirements.

Example: *An organization already ISO/IEC 27001-certified can reuse its existing access control and incident response evidence for the AI security checklist in Section 10, adding only the AI-specific gaps — reducing duplicate audit work significantly.*

10. AI Security & Privacy Checklist

Data classification

Confirm the sensitivity level of data an AI system touches (public, internal, confidential, restricted) and ensure the system's controls match that classification.

Access controls

Restrict who can configure, query, or export data from the AI system; apply least-privilege and role-based access.

Prompt / data protection

Guard against prompt injection and data leakage through system prompts; avoid embedding secrets or sensitive data directly in prompts.

Third-party AI tools

Maintain an approved list of external AI tools; block or flag use of unapproved tools where feasible.

Model security

Assess exposure to model theft, adversarial inputs, and manipulation of outputs; apply monitoring for anomalous use patterns.

Data retention

Define how long prompts, outputs, and training data are retained, and ensure vendor retention practices are documented and acceptable.

Sensitive information exposure

Test whether the system can be prompted into revealing sensitive internal or personal data it was not intended to expose.

AI incident response

Extend existing security incident processes to cover AI-specific scenarios: data leakage via prompts, biased or harmful outputs, and model manipulation.

11. AI Vendor & Third-Party Assessment

Ask these questions before signing with any AI vendor. Keep the answers on file — they are often required for audits and regulatory inquiries.

1. Where is data processed? — Identify the physical/jurisdictional location of processing, especially for cross-border data transfer rules.

2. Is customer data used for model training? — Confirm whether your data trains the vendor's shared models, or is used only to serve your own requests.
3. What security controls are available? — Encryption at rest/in transit, access logging, SOC 2 or equivalent certifications.
4. How is data retained? — Retention periods for inputs, outputs, and logs, and whether you can request deletion.
5. What certifications does the vendor hold? — SOC 2, ISO/IEC 27001, ISO/IEC 42001, or sector-specific certifications.
6. How are AI risks disclosed? — Does the vendor publish model cards, known limitations, or bias testing results?
7. What happens when the AI system fails? — SLAs, fallback procedures, and who is liable for AI-driven errors.

Example: *A vendor questionnaire response stating 'customer prompts may be used to improve our models unless you opt out in the enterprise settings' is a material finding — it should be flagged, opted out where required, and documented in the vendor assessment file regardless of how minor it seems.*

12. AI Governance Audit Checklist

Use this checklist to prepare for or conduct an internal audit of your AI governance program.

☐ Governance structure — roles, committee, and escalation paths are documented and active.

☐ AI inventory — complete, current, and reviewed within the last quarter.

☐ Risk assessments — completed for every system in the inventory, with scores on file.

☐ Policies — the AI use policy is published, current, and acknowledged by staff.

☐ Documentation — risk assessments, vendor questionnaires, and approvals are retained and retrievable.

☐ Security controls — access controls and monitoring match each system's risk level.

☐ Privacy controls — data minimization and lawful basis are documented for systems using personal data.

☐ Human oversight — override and escalation mechanisms exist and have been tested.

☐ Monitoring — a defined cadence exists for re-reviewing systems and risks.

☐ Incident management — AI incidents are logged, and the process has been exercised or tested.

☐ Third-party risk — vendor assessments are complete and current for all in-use AI vendors.

13. AI Governance Maturity Model

Use this model to assess where your organization stands today and to set a realistic target for the next 12 months.

Level	Name	Description
1	Initial	AI adoption is largely unmanaged; tools are adopted ad hoc with no inventory, ownership, or risk review.
2	Developing	Basic policies and ownership are emerging; some systems are documented, but coverage is inconsistent.
3	Defined	Standardized governance processes are established; inventory, risk assessment, and policy are applied consistently to new AI systems.
4	Managed	AI risks and controls are actively monitored; metrics are tracked, and reviews happen on a defined cadence.
5	Optimized	Governance is integrated into enterprise strategy and continuously improved; AI risk data informs broader business and technology decisions.

Example: *Most organizations starting from scratch are at Level 1 or 2. A realistic 12-month target is Level 3 — with a documented inventory, consistent risk scoring, and a published policy — before pursuing the ongoing monitoring maturity of Level 4.*

14. 30-Day AI Governance Action Plan

A practical, week-by-week plan to stand up the fundamentals of a governance program in one month.

Week 1 — Discover

- Identify AI systems currently in use across departments.
- Identify stakeholders who should be involved in governance (IT, legal, HR, business owners).
- Document current practices — however informal — for how AI tools are adopted today.

Week 2 — Assess

- Conduct risk assessments for the highest-visibility or highest-risk systems first.
- Identify compliance gaps against applicable regulations and standards.
- Review third-party AI tools already in use for vendor risk.

Week 3 — Govern

- Assign accountability using the roles matrix from Section 4.
- Create or adapt policies using the starter template from Section 8.
- Define controls proportional to each system's risk level.

Week 4 — Monitor

- Establish KPIs using the metrics dashboard in Section 15.
- Set a review cycle (recommended: quarterly) for the inventory and risk register.
- Plan the first internal audit and continuous improvement checkpoint.

Example: *A 200-person software company used this plan to go from zero formal governance to a published policy, a 40-item inventory, and named owners for every AI system in exactly 30 days — proving that speed and rigor aren't mutually exclusive when the framework is already laid out.*

15. AI Governance Metrics Dashboard

Track these KPIs monthly or quarterly and report them to leadership to show the program is working — and where it needs support.

KPI	Why It Matters
% of AI systems inventoried	Shows visibility — the foundation everything else depends on.
% of AI systems risk-assessed	Shows how much of the inventory has moved past discovery into evaluation.
% with assigned owners	Indicates whether accountability gaps remain.

% meeting governance requirements	The core measure of program effectiveness.
Number of AI incidents	Tracks real-world impact and trend over time.
Number of unresolved AI risks	Highlights backlog and prioritization needs.
Third-party AI assessments completed	Tracks vendor risk coverage.
Employee AI governance training completion	Measures how well the policy is understood and followed day-to-day.

16. AI Governance Resources

AI governance frameworks

- NIST AI Risk Management Framework (AI RMF)
- ISO/IEC 42001 — AI Management System

Standards

- ISO/IEC 27001 — Information Security Management
- ISO/IEC 23894 — AI Risk Management guidance

Regulatory resources

- EU AI Act official text and guidance
- Sector-specific regulators' AI guidance (e.g., financial services, healthcare)

Risk assessment resources

- NIST AI RMF Playbook
- Algorithmic impact assessment templates from public-sector guidance bodies

Responsible AI references

- OECD AI Principles
- Published responsible-AI frameworks from major cloud/AI providers

Recommended governance tools

- AI inventory and risk-register software (spreadsheet-based to start; dedicated GRC platforms as you scale)
- Model monitoring tools for tracking accuracy and drift in production

17. Final Checklist: Are You Ready for AI Governance?

Print this page and use it as a quick self-assessment.

☐ AI systems identified

☐ AI owners assigned

☐ AI risks assessed

☐ AI policy established

☐ Data/privacy controls defined

☐ Security controls implemented

☐ Human oversight established

☐ Vendors assessed

☐ Monitoring process established

☐ Audit process defined

If you checked fewer than seven boxes, start with Section 14's 30-Day Action Plan.

18. Next Step: Build Your AI Governance Expertise

Ready to Take AI Governance Further?

Building the templates in this kit is the first step. Running a governance program well over time — across a growing AI footprint, shifting regulation, and cross-functional stakeholders — takes dedicated expertise.

Organizations need professionals who can connect AI governance, risk management, compliance, security, and responsible AI into a practical governance strategy that holds up under audit and scales with adoption.

GSDC Certified AI GRC Professional Certification can help professionals develop the knowledge and skills needed to manage AI governance and GRC initiatives — turning the templates in this kit into a program that lasts.

GSDC
Global Skill Development Council

CERTIFIED AI GRC PROFESSIONAL

THE AI GOVERNANCE, RISK & COMPLIANCE (AI GRC) CERTIFICATION PROGRAM IS GLOBALLY DESIGNED TO STRENGTHEN AI GOVERNANCE, RISK MANAGEMENT, REGULATORY COMPLIANCE, AND RESPONSIBLE AI IMPLEMENTATION SKILLS, ENABLING PROFESSIONALS TO ESTABLISH TRUSTWORTHY, SECURE, AND COMPLIANT AI SYSTEMS WHILE ALIGNING INNOVATION WITH ORGANIZATIONAL AND REGULATORY REQUIREMENTS.

ABOUT GSDC CERTIFICATION

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.

LEARNING OBJECTIVE

- Prove your knowledge of responsible AI principles including fairness, transparency, explainability, and human oversight in AI systems
- Understand how to design and implement AI security controls, manage adversarial risks, and build trustworthy AI programs across the organization

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now

www.gsdcouncil.org