GSDC
Global Skill Development Council

Globally Recognised Certification Body
190+ countries · 2,50,000+ Certified Professionals

AI Governance Checklist for Organizations

www.gsdcouncil.org

How to Use This Checklist

This checklist is organised into 14 governance domains, each representing a critical pillar of a robust AI governance programme. For each domain, you will find a set of actionable items to evaluate. Work through each item and mark it as complete, in progress, or a gap requiring urgent attention.

Use the checklist as a diagnostic tool during initial programme setup, as part of a periodic governance review cycle, or in preparation for a regulatory audit. The checklist is designed to be modular — you may prioritise domains based on your organisation's risk profile, regulatory environment, and current AI maturity.

Each section includes context to help teams understand why each domain matters, not just what to check. Share the completed checklist with senior leadership and your board to drive governance investment decisions.

Scoring Guide

01

Complete

Control is in place, documented, and operating effectively.

02

In Progress

Work has begun but the control is not yet fully operational.

03

Gap

No control exists or significant deficiencies have been identified.

DOMAIN 1

AI Governance Structure

Effective AI governance begins at the top. Without a clearly defined governance structure, accountability is diffuse, decisions are slow, and oversight is inconsistent. Your governance structure should establish who is responsible for AI oversight, how decisions are made, and how escalation paths function across the organisation. This is the scaffolding upon which every other governance domain depends.

- [ ] An AI governance committee or board-level oversight body has been formally established with defined terms of reference.
- [ ] Executive sponsorship for AI governance is assigned, with named accountability at C-suite level.
- [ ] The governance structure spans relevant functions: legal, compliance, IT, risk, HR, and business units.
- [ ] Decision-making authority for AI deployment, modification, and decommissioning is clearly documented.
- [ ] Escalation pathways for AI-related issues, risks, and incidents are defined and tested.
- [ ] Governance meetings are scheduled at regular cadence with documented minutes and actions.

- [ ] The AI governance structure is reviewed annually and updated to reflect organisational or regulatory changes.
- [ ] Board or audit committee receives regular AI governance reporting.

DOMAIN 2

AI Inventory

You cannot govern what you cannot see. A comprehensive AI inventory — sometimes called an AI register — is a foundational requirement for any mature governance programme. It provides visibility across all AI systems in use, enabling risk assessment, policy application, and audit readiness. Without it, organisations operate blind to the full extent of their AI exposure.

Inventory Essentials

- [ ] A centralised AI inventory or register is maintained listing all AI systems in use, including those embedded in third-party software.
- [ ] Each entry records the system name, vendor (if applicable), business owner, intended use case, and deployment environment.
- [ ] The risk classification of each AI system is documented (e.g., high, medium, low risk).
- [ ] Date of deployment, last review date, and next scheduled review are recorded for each system.

Inventory Maintenance

- [ ] A process exists to capture newly deployed AI systems within a defined timeframe of go-live.
- [ ] The inventory is reviewed and reconciled at least quarterly against procurement records and IT asset lists.
- [ ] Retired or decommissioned AI systems are archived with a record of decommissioning rationale and date.
- [ ] The inventory is accessible to governance, risk, compliance, and audit functions on demand.

A well-maintained AI inventory is often the first document requested during a regulatory audit or internal governance review. Treat it as a living record, not a one-time exercise.

DOMAIN 3

Roles and Responsibilities

Clear role definition eliminates the ambiguity that leads to governance failures. Every AI system should have a named owner, and every governance function should have a defined remit. Overlapping or unclear responsibilities are among the most common root causes of AI risk incidents. A well-designed RACI matrix for AI governance translates accountability from policy into practice.

AI System Owner

- [ ] Every AI system has a named business owner accountable for its performance and governance compliance.
- [ ] The AI system owner's responsibilities are formally documented and accepted in writing.

AI Risk & Compliance

- [ ] A named individual or team holds responsibility for AI risk assessment and policy compliance monitoring.
- [ ] Compliance roles have formal authority to halt or escalate AI deployments that breach policy.

AI Developers & Data Scientists

- [ ] Developers and data scientists are aware of their governance obligations at each stage of the AI lifecycle.
- [ ] Responsible AI principles are embedded in development standards and onboarding.

Legal & Ethics

- [ ] Legal counsel is engaged for AI use cases with significant regulatory or liability exposure.
- [ ] An ethics advisory function or process exists for novel or high-risk AI applications.

DOMAIN 4

AI Risk Assessment

Risk assessment is the engine of AI governance. Without a structured, repeatable process for evaluating AI risk, organisations cannot make informed decisions about which systems to deploy, what controls to apply, or where to invest governance resources. AI risk is multidimensional — it spans technical failure, bias and fairness, legal and regulatory exposure, reputational harm, and operational disruption.

Risk assessments should be conducted before deployment, at significant change events, and on a periodic basis throughout the AI system lifecycle. The depth and formality of the assessment should be proportionate to the system's risk classification.

Pre-Deployment Assessment

- [ ] A structured AI risk assessment is completed before any new AI system is deployed into production.
- [ ] Risk assessments consider bias, fairness, transparency, security, privacy, and regulatory compliance dimensions.
- [ ] Risk classification (high/medium/low) is assigned and documented with supporting rationale.

Ongoing Assessment

- [ ] Periodic risk reassessments are scheduled and completed for all active AI systems.
- [ ] Material changes to AI systems (model updates, new data inputs, expanded use cases) trigger a reassessment.
- [ ] Risk assessment findings are tracked to resolution with assigned owners and target dates.

High-Risk AI

- [ ] High-risk AI systems are subject to enhanced due diligence, including independent review or external audit.
- [ ] A formal impact assessment (e.g., Data Protection Impact Assessment, Algorithmic Impact Assessment) is conducted where required by regulation or policy.

DOMAIN 5

AI Policies

Policies are the formal expression of an organisation's expectations, standards, and boundaries for AI use. A well-structured policy framework provides clarity for all stakeholders — from developers and business users to vendors and auditors. Without documented policies, governance is informal, inconsistent, and impossible to enforce or audit.

Policies should cover the full AI lifecycle and address the organisation's specific risk profile, sector regulations, and ethical commitments. They must be reviewed regularly and kept current with evolving legal and technical landscapes.

Acceptable Use Policy

- [ ] An AI acceptable use policy is in place defining permitted and prohibited AI applications.
- [ ] The policy addresses employee use of consumer AI tools and generative AI platforms.

AI Ethics Policy

- [ ] An AI ethics policy articulates the organisation's principles on fairness, accountability, transparency, and non-maleficence.
- [ ] Ethics principles are operationalised with practical guidance, not limited to aspirational statements.

AI Procurement Policy

- [ ] A policy governs the procurement and onboarding of AI systems and vendors, including governance requirements.
- [ ] Procurement policy requires vendor disclosure of AI capabilities embedded in software products.

Policy Governance

- [ ] All AI policies are version-controlled, owner-assigned, and subject to at least annual review.
- [ ] Policies are accessible to all relevant staff and training is provided to support comprehension and compliance.

DOMAIN 6

Data Governance

AI systems are only as trustworthy as the data that trains and drives them. Data governance and AI governance are inseparable disciplines. Poor data quality, insufficient data lineage, unaddressed privacy risks, and inadequate consent management can all undermine AI system integrity, produce biased outputs, and expose the organisation to significant regulatory liability — particularly under data protection frameworks such as the GDPR or the EU AI Act.

Data Quality & Lineage

- [ ] Training data sources are documented, including origin, collection method, and any known limitations or biases.
- [ ] Data quality standards are defined and applied to datasets used in AI model training and inference.
- [ ] Data lineage is traceable from source to model output for all high-risk AI systems.
- [ ] Data drift monitoring is in place to detect changes in input data distributions that may affect model performance.

Privacy & Compliance

- [ ] Personal data used in AI systems is identified, lawful basis for processing is established, and DPIA conducted where required.
- [ ] Data minimisation principles are applied — AI systems use only the data necessary for the intended purpose.
- [ ] Data retention and deletion obligations are met for AI training data and inference logs.
- [ ] Cross-border data transfer restrictions are assessed and addressed for AI systems using cloud or third-party infrastructure.

- [ ] Data governance accountabilities for AI-related datasets are assigned to named data owners.
- [ ] Data governance policies explicitly address AI use cases and are aligned with the organisation's AI governance framework.

DOMAIN 7

Transparency and Explainability

Transparency and explainability are central to accountable AI. Stakeholders — whether customers, employees, regulators, or auditors — have a legitimate interest in understanding how AI systems make decisions, particularly when those decisions have a material impact on individuals or organisational outcomes. Explainability is not merely a technical feature; it is a governance obligation in an increasing number of regulatory contexts.

Organisations must balance the need for explainability with practical constraints such as model complexity, commercial confidentiality, and computational cost. A risk-proportionate approach — requiring greater explainability for higher-risk applications — is a widely accepted standard.

→ Explainability Standards

- [ ] Explainability requirements are defined for each AI system commensurate with its risk classification and the nature of decisions it informs.
- [ ] Explainability techniques (e.g., SHAP, LIME, model cards) are applied to high-risk AI systems and outputs documented.

→ User-Facing Disclosure

- [ ] Individuals subject to AI-driven decisions are informed that AI is being used and, where required, given a right to explanation or human review.
- [ ] Customer-facing AI interactions (e.g., chatbots, recommendation engines) are clearly identified as AI-powered.

→ Internal Transparency

- [ ] Business users of AI systems understand the basis for AI outputs and are not over-reliant on black-box recommendations.
- [ ] Model cards or system documentation describe intended use, limitations, and performance characteristics for internal audiences.

DOMAIN 8

Human Oversight

Human oversight is a foundational safeguard in responsible AI deployment. No matter how capable or well-tested an AI system is, meaningful human control must be maintained — particularly for decisions with significant consequences for individuals, the organisation, or society. The EU AI Act, the NIST AI Risk Management Framework, and other leading governance standards all underscore human oversight as a non-negotiable requirement for high-risk AI.

Human oversight is not simply about having a person in the loop. It requires that the human reviewer has sufficient understanding, authority, and time to make a genuine independent judgement — not merely to rubber-stamp an AI recommendation.

1

Meaningful Review

- [ ] Human reviewers of AI-assisted decisions are trained to understand the AI system's outputs, limitations, and potential failure modes.
- [ ] Review processes are designed to prevent automation bias — where humans accept AI recommendations without genuine scrutiny.

2

Override Capability

- [ ] A clearly defined and technically enabled mechanism exists for humans to override, reject, or modify AI recommendations at any stage.
- [ ] Override decisions are logged, with rationale, to support audit and continuous improvement.

3

High-Stakes Decisions

- [ ] AI systems are not used to make final, automated decisions in high-stakes contexts (e.g., credit, employment, healthcare) without human review, unless explicitly permitted by law.
- [ ] Escalation routes exist for AI outputs that fall outside expected parameters or confidence thresholds.

4

Oversight Audit Trail

- [ ] Human oversight activities are logged with sufficient detail to demonstrate that review was substantive, not nominal.
- [ ] Oversight metrics (e.g., override rates, review times) are tracked and reported to governance bodies.

DOMAIN 9

AI Security

AI systems introduce a distinct and expanding set of security vulnerabilities that differ materially from traditional software risks. Adversarial attacks, model inversion, data poisoning, prompt injection, and model theft are not hypothetical threats — they are documented, real-world attack vectors that organisations must actively address. A mature AI security programme integrates AI-specific controls into the broader information security management framework.

Security testing for AI systems should begin at the development stage and continue throughout the operational lifecycle. Security requirements should be embedded in AI procurement and vendor management processes, ensuring that third-party AI systems meet the organisation's security standards.

Threat Modelling

- [ ] AI-specific threat modelling is conducted for high-risk AI systems prior to deployment, covering adversarial attacks, data poisoning, and model extraction.
- [ ] Threat models are reviewed when significant changes are made to the AI system or its operating environment.

Access Controls

- [ ] Access to AI models, training data, and inference infrastructure is restricted to authorised personnel with role-appropriate permissions.
- [ ] Privileged access to AI systems is subject to enhanced controls, logging, and periodic review.

Secure Development

- [ ] AI systems are developed under a secure development lifecycle (SDL) that includes AI-specific security testing (e.g., adversarial robustness testing, red-teaming).
- [ ] Prompt injection and jailbreak risks are assessed and mitigated for generative AI applications.

Incident Readiness

- [ ] AI-specific security incidents are addressed within the organisation's broader incident response plan.
- [ ] Procedures exist for isolating and disabling AI systems in the event of a confirmed security compromise.

DOMAIN 10

Third-Party AI Systems

The majority of AI used by organisations today is not built in-house — it is procured from vendors, embedded in SaaS platforms, accessed via APIs, or delivered as part of managed services. This creates significant governance complexity. The organisation remains accountable for the AI's impact on its customers, employees, and operations, even when it does not own or control the underlying model.

Third-party AI governance requires proactive due diligence at procurement, contractual protections, ongoing monitoring, and clear exit strategies. Regulators increasingly expect organisations to demonstrate control over their AI supply chain.

Procurement Due Diligence

- [ ] A due diligence process is applied to AI vendors and systems prior to procurement, covering risk, security, data practices, and regulatory compliance.
- [ ] Vendors are required to disclose what AI capabilities are embedded in their products and how they operate.
- [ ] AI vendor assessments include review of the vendor's own AI governance programme and responsible AI commitments.
- [ ] High-risk third-party AI systems are subject to independent technical review or third-party audit before deployment.

Contractual & Ongoing Controls

- [ ] Contracts with AI vendors include provisions for transparency, audit rights, incident notification, and data protection obligations.
- [ ] SLAs include AI-specific performance, explainability, and availability requirements.
- [ ] Third-party AI systems are included in the organisation's AI inventory and risk assessment process.
- [ ] Vendor AI governance performance is reviewed at least annually as part of supplier management.

- [ ] Exit strategies are documented for critical third-party AI dependencies, including data portability and operational continuity plans.

DOMAIN 11

Shadow AI

Shadow AI — the use of AI tools and systems outside of organisational oversight, approval, or knowledge — is one of the fastest-growing governance risks facing organisations today. Employees using consumer generative AI tools, installing AI-powered browser extensions, or feeding sensitive data into unapproved platforms represent a significant and often invisible risk to data security, regulatory compliance, and organisational reputation.

Addressing shadow AI requires a combination of clear policy, effective awareness training, technical controls, and a culture that channels AI enthusiasm into governed, approved channels rather than suppressing it entirely. Organisations that over-restrict without providing approved alternatives often drive shadow AI underground rather than eliminating it.

Detection & Visibility

- [ ] Technical controls (e.g., network monitoring, endpoint management, DLP tools) are in place to detect the use of unapproved AI tools on organisational systems.
- [ ] A process exists to identify and assess newly discovered shadow AI usage, with defined escalation and remediation steps.
- [ ] Periodic audits of software and application usage include a review for unapproved AI tools.

Prevention & Culture

- [ ] Employees are aware of the acceptable use policy for AI tools and understand specifically which tools are permitted and prohibited.
- [ ] A simple, low-friction process exists for employees to request approval for new AI tools, reducing incentives for informal use.
- [ ] Training on shadow AI risks is included in general AI literacy and security awareness programmes.

- [ ] Approved alternatives to common consumer AI tools are provided and communicated to employees, reducing the demand for shadow AI.

DOMAIN 12

AI Incident Management

AI systems can and do fail — and when they do, the consequences can be swift, severe, and highly public. An AI incident management capability ensures that the organisation can detect, respond to, and learn from AI-related failures, harms, and near misses in a structured and timely manner. Incident management for AI is distinct from general IT incident management in important respects: it must address harms to individuals and groups, bias and fairness failures, regulatory breach scenarios, and reputational crises, in addition to technical outages.

1

Detect

Monitoring and reporting mechanisms identify potential AI incidents, including automated alerts and human escalation routes.

2

Assess

Incidents are triaged and classified by severity, nature, and potential regulatory or reputational impact within defined timeframes.

3

Respond

Defined response playbooks activate the appropriate team, communications, and remediation actions for each incident type.

4

Learn

Post-incident reviews identify root causes, governance gaps, and improvements to prevent recurrence.

- [ ] An AI incident response plan is documented and tested at least annually, including AI-specific scenarios.
- [ ] Regulatory notification obligations for AI incidents (e.g., data breaches caused by AI) are understood and built into the response plan.
- [ ] A log of all AI incidents, near misses, and complaints is maintained and reviewed by the governance committee.
- [ ] Affected individuals and regulators are notified in line with legal obligations and within required timeframes.

DOMAIN 13

Monitoring and Continuous Improvement

AI governance is not a one-time exercise. AI systems evolve, data distributions shift, regulatory requirements change, and organisational risk appetites develop over time. A robust monitoring and continuous improvement programme ensures that AI systems remain aligned with governance requirements throughout their operational life — and that the governance framework itself is regularly reviewed and enhanced.

Continuous improvement in AI governance requires feedback loops from operations, incidents, audits, and external developments. Governance bodies should regularly review performance data, risk indicators, and emerging best practices to maintain a governance programme that is current, effective, and proportionate.

Monitor Performance

- [ ] KPIs and monitoring metrics are defined for each AI system, covering accuracy, fairness, availability, and governance compliance.

Review & Audit

- [ ] Periodic governance reviews assess AI system performance, risk control effectiveness, and policy compliance across the AI portfolio.

Implement & Update

- [ ] Improvements are implemented and their effectiveness verified before closing. Governance policies and registers are updated accordingly.

Identify Improvements

- [ ] Findings from monitoring, audits, and incidents are translated into specific, time-bound improvement actions with assigned owners.

- [ ] External developments — regulatory changes, new standards, peer incidents — are monitored and incorporated into the governance improvement process.
- [ ] Governance programme maturity is assessed annually against recognised frameworks (e.g., NIST AI RMF, ISO/IEC 42001).

DOMAIN 14

Documentation and Audit Readiness

Documentation is the foundation of accountability. Regulators, auditors, and courts will expect organisations to demonstrate — not merely claim — that AI governance controls are in place and operating effectively. In the absence of robust documentation, even organisations with mature governance practices will struggle to evidence compliance. Audit readiness is not a reactive posture taken in the days before an inspection; it is a discipline maintained continuously throughout the governance lifecycle.

Good documentation practices serve a dual purpose: they hold the organisation accountable to its own standards and they provide the evidence base needed to demonstrate compliance to external parties. Documentation should be clear, proportionate, version-controlled, and retained for periods appropriate to the regulatory context.

Core Documentation

- [ ] AI system documentation is maintained for each system in the inventory, covering purpose, design, training data, performance benchmarks, and known limitations.
- [ ] Risk assessments, impact assessments, and their outcomes are documented and version-controlled.
- [ ] Governance committee minutes, decisions, and escalations are formally recorded and retained.
- [ ] Policies, procedures, and controls are documented with version history and approved owner sign-off.

Audit Trail

- [ ] AI system logs sufficient to reconstruct decision pathways are retained in line with legal and regulatory requirements.
- [ ] Training data records, model version histories, and change logs are maintained and accessible.
- [ ] Human oversight activities, including override decisions and review records, are logged with appropriate detail.
- [ ] Incident records, post-incident reviews, and corrective action plans are retained and accessible to auditors.

- [ ] A documentation retention schedule aligned to applicable regulatory requirements is defined and implemented for all AI-related records.
- [ ] Documentation is stored in a secure, searchable repository with access controls and backup procedures.

AI Risk Classification Quick Reference

Use this reference guide to help classify AI systems within your inventory. Risk classification should be applied consistently and reviewed whenever an AI system's use case, data inputs, or operational context changes significantly. Classification drives the depth of governance controls applied to each system.

Risk Level	Typical Characteristics	Example Use Cases	Governance Intensity
Critical / High	Automated decisions affecting individuals' rights, safety, or significant interests; high regulatory exposure.	Credit scoring, HR screening, medical diagnosis support, law enforcement tools.	Full lifecycle governance, independent audit, mandatory human oversight, impact assessment required.
Medium	Decisions with material business impact; moderate regulatory risk; significant operational dependency.	Customer churn prediction, dynamic pricing, fraud detection, content moderation.	Standard governance controls, periodic risk assessment, documented oversight, explainability requirements.
Low	Decisions with limited individual impact; low regulatory exposure; easily reversible outcomes.	Internal productivity tools, document summarisation, meeting transcription, scheduling assistants.	Inventory registration, acceptable use policy compliance, basic monitoring, periodic review.
Minimal	No individual impact; purely internal analytics or process support with no consequential decisions.	Internal search, code completion tools, background automation with human review of all outputs.	Acceptable use policy compliance, inclusion in inventory, no additional controls required.

AI Governance Maturity Overview

Understanding your organisation's current AI governance maturity is the first step towards targeted, effective improvement. Use this infographic to locate your organisation across the five maturity levels. Most organisations beginning a formal AI governance journey fall between Level 1 and Level 2. The goal is not to reach Level 5 in every domain simultaneously, but to progress systematically, prioritising domains that present the greatest risk exposure.

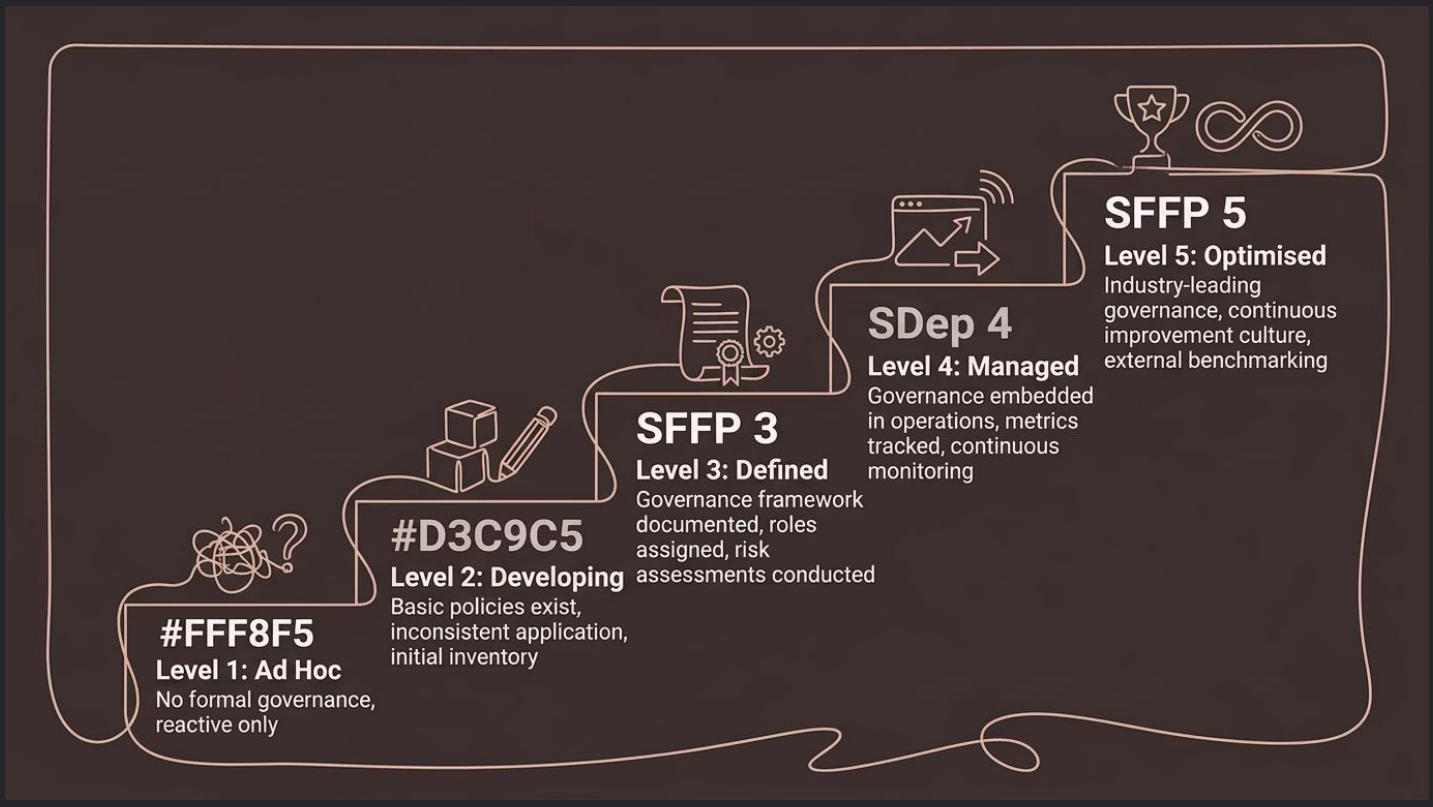

Progress across maturity levels requires commitment from leadership, sustained resourcing, and integration of governance into the AI development and procurement lifecycle — not as a bolt-on, but as a core organisational capability.

Governance Review Cadence

Effective AI governance requires structured, recurring review activities at multiple levels of the organisation. Ad hoc reviews are insufficient — governance must be embedded as a standing operational discipline. The table below outlines a recommended review cadence that organisations can adapt to their own governance structure, risk profile, and regulatory context.

Activity	Frequency	Responsible Party	Output
AI inventory reconciliation	Quarterly	AI Programme Manager / IT	Updated AI register
High-risk AI system review	Quarterly	AI System Owner / Risk	Risk assessment update, performance report
AI governance committee meeting	Quarterly	Governance Committee Chair	Minutes, decisions, action log
Policy review	Annual	Compliance / Legal	Updated policies, version log
Full AI governance audit	Annual	Internal Audit / External Auditor	Audit report, management responses
Maturity assessment	Annual	AI Programme Manager	Maturity scorecard, improvement roadmap
Board / Audit Committee briefing	Semi-annual	CRO / CISO / AI Lead	Governance dashboard, risk summary
Incident response plan test	Annual	Risk / Security / Operations	Test report, lessons learned

Next Steps: Closing Your Governance Gaps

Completing this checklist is the beginning of your AI governance journey, not the end. The real value lies in translating the gaps you have identified into a prioritised, resourced, and time-bound improvement plan. Use your findings to brief senior leadership, secure investment, and build the cross-functional coalitions necessary to drive governance maturity across your organisation.

1 Prioritise by Risk

Rank identified gaps by the risk they present. Focus initial effort on the highest-risk AI systems and the domains with the greatest governance deficiencies. Do not attempt to address every gap simultaneously — prioritisation is essential.

2 Assign Ownership

Every gap should have a named owner, a target completion date, and defined success criteria. Unowned actions do not get completed. Build accountability into your improvement plan from the outset.

3 Engage Leadership

Present your governance gap analysis to your board or executive committee. AI governance requires top-down commitment. Leadership engagement unlocks the resources, authority, and organisational attention needed to make meaningful progress.

4 Benchmark and Iterate

Align your governance programme to recognised frameworks such as the NIST AI Risk Management Framework, ISO/IEC 42001, or the EU AI Act requirements. Repeat this checklist assessment annually to track progress and identify emerging gaps as your AI landscape evolves.

AI governance is a journey of continuous improvement. Organisations that build governance maturity systematically — one domain at a time — will be better positioned to deploy AI responsibly, maintain stakeholder trust, and meet the demands of an increasingly regulated AI environment.

GSDC
Global Skill Development Council

Globally Recognised Certification Body
190+ countries · 2,50,000+ Certified Professionals

IAF MEMBER OF MULTILATERAL RECOGNITION ARRANGEMENT
IAS ACCREDITED Personnel Certification Body

CERTIFIED AI GOVERNANCE PROFESSIONAL

ABOUT GSDC CERTIFICATION

EBOOK

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

LEARNING MATERIALS

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

CREATED BY EXPERTS

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

LEARNING OBJECTIVE

- Gain a strong understanding of AI governance principles and practices
- Apply AI governance concepts to real-world organizational needs
- Develop skills to identify and manage AI-related risks
- Build effective AI governance practices for responsible AI adoption

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

Enroll Now

www.gsdcouncil.org