AI GOVERNANCE GUIDE

AI Governance: Best Practices

General best practices for the teams that govern AI, built in maturity stages, for any organization

Roger Estrada · AI Governance · AI Strategy · 2026

AI governance decides what AI is allowed to do, who is accountable, and how leadership knows it is still operating within limits. It works only when the teams that own risk, compliance, legal, data, security, vendors, people and the business act together. This document describes the role of each team and the practices that mature it. It applies to any industry and any size; regulated sectors add their own rules.

> **THE CORE IDEA**
>
> Governance is an operating discipline with named owners, risk tiers, evidence and stop authority. A policy that is approved once and filed is not governance. Govern the system that is running today, not the one that was approved last year.

Guiding Principles

- **Accountable.** Every AI system has a named owner and an executive sponsor.
- **Proportionate.** Controls scale with the consequence of a wrong output and the authority the system holds.
- **Transparent.** The organization can explain what a system does, what data it uses and how decisions were reached.
- **Fair.** Outcomes are tested for harmful bias, and people can seek review by a human.
- **Secure and private.** Data access follows least privilege and existing data classification.
- **Reliable.** Systems are tested before release and monitored after.
- **Human-centered.** People keep meaningful authority over consequential decisions.

Governance Maturity

Find the stage that describes your organization today, then adopt the practices for the next stage. Each stage builds on the one before.

Stage	What it looks like	What moves you up
1. Ad hoc	Scattered tools, no inventory, rules left to individuals	Sponsor, council, inventory, acceptable-use policy
2. Defined	Policies, tiers and owners exist; reviews happen at approval	Monitoring, re-approval triggers, vendor controls, training
3. Managed	Systems monitored; evidence collected; leadership sees regular reporting	Independent testing, incident exercises, reuse of controls, value tracking

Stage	What it looks like	What moves you up
4. Optimized	Continuous assurance, fast approved paths, governance adapts to new technology and rules	Yearly refresh, horizon scanning, lessons learned

Best Practices by Team

Stage 1 is the starting point with no formal practice, so the table begins at Stage 2.

Team	Stage 2: Defined	Stage 3: Managed	Stage 4: Optimized
Executive sponsor and board Sets direction and accepts risk	Appoint the sponsor and charter a cross-functional council. Approve the AI policy and risk appetite. Receive an annual briefing.	Receive quarterly reports on inventory, tiers, incidents and value. Decide exceptions. Ask the five oversight questions below.	Set AI strategy and funding by value and risk. Review emerging risks yearly. Require independent assurance.
Risk management Sets tiers and owns the risk register	Create tiering criteria and authority levels. Keep an AI risk register with a named acceptor for each residual risk. Include AI in the enterprise risk framework.	Challenge and validate high-tier systems. Track risk indicators and trigger events. Set limits on residual risk.	Express AI risk in appetite metrics. Run scenario exercises. Use risk data in funding decisions.
Compliance Maps use to law, standards and policy	Map each use case to applicable laws, standards and contract terms. Set disclosure, fairness-testing and complaint rules.	Run bias and explainability tests on schedule. Monitor complaints tied to AI. Keep evidence current for regulators.	Own regulatory tracking. Automate control mapping and evidence collection. Engage regulators early on new uses.
Legal and privacy Protects the organization in contracts and data use	Set contract terms for data use, model training, IP, liability, audit, deletion and exit. Define retention for prompts, outputs and logs. Run privacy assessments.	Review high-tier and external-facing uses. Keep records and discovery ready. Negotiate vendor changes.	Maintain standard AI clauses by vendor type. Track legislation and litigation. Advise on agents and automated decisions.
Data governance Controls what AI can see and use	Apply data classification to AI. Publish which data may enter which tools. Document sources, permitted purposes and retention.	Check quality, lineage and freshness for decision-support uses. Screen data for bias. Test that retrieval respects permissions.	Automate lineage and quality monitoring. Maintain governed, reusable data for AI. Measure data readiness.
Information security Protects systems and data	Give each AI system its own identity with least privilege. Approve connectors. Log actions. Define an AI incident.	Red-team high-risk systems for prompt injection and data leakage. Monitor for misuse. Test the stop procedure.	Test agents continuously. Feed AI events into security operations. Secure the AI supply chain.
Vendor management Governs third-party and embedded AI	Add AI questions to due diligence. Require change notice, audit and exit rights. Inventory vendor-embedded AI.	Reassess vendors on material releases. Watch model and subprocessor changes. Track provider concentration.	Tier vendors by AI dependence. Seek shared assurance. Keep tested exit plans.

Team	Stage 2: Defined	Stage 3: Managed	Stage 4: Optimized
HR and learning Prepares people	Publish acceptable-use rules. Deliver role-based training. Open a reporting route for concerns.	Track completion and understanding. Update training after incidents. Support role changes.	Build AI literacy into hiring, objectives and leadership development. Measure a culture that challenges output.
Business owners Own purpose, outcome and value	Own the purpose, workflow and outcome. Fix the process before automating. Give reviewers the evidence and time to challenge output.	Monitor performance, overrides and fallback readiness. Track value for each use case.	Reuse approved patterns. Redesign work around outcomes. Retire what no longer delivers.
Internal audit Provides independent assurance	Add AI to the audit universe. Review the inventory and policies.	Test whether controls operate as described. Sample human review quality.	Audit continuously from automated evidence. Advise on emerging risks.

Practices Every Team Shares

- **One inventory.** Record every AI system, whether built, bought or embedded in a vendor product, plus employee-adopted tools such as ChatGPT, Claude and Gemini. Capture purpose, data, permissions, owner and tier. You cannot govern what you cannot see.
- **Risk tiers.** Classify by consequence and by what the system can do, and set controls by tier.
- **One framework.** Use the NIST AI Risk Management Framework as the backbone, with ISO/IEC 42001 as an option. Govern covers policy, accountability and culture; Map covers context and impact; Measure covers testing and monitoring; Manage covers treatment, response and vendors.
- One intake and one decision record. A requester submits once. Each team reviews in parallel against pass, discuss and block criteria it agreed in advance, and each owns the judgment in its field.
- **Real human oversight.** A person who clicks approve on pre-ranked output provides little oversight. Define what the reviewer is responsible for, show the evidence, allow time, and sample reviews to confirm they are substantive.
- **Re-approval triggers.** Send a system back through review on a new model or vendor release, expanded data or permissions, a move from recommending to acting, an unapproved new use, or evidence that review has become a formality.
- **An incident plan.** Cover containment, notification, recovery and a named person who can stop a system.

Tier	Examples	Minimum controls
Low	Drafting, brainstorming, notes on non-sensitive data	Approved tool, acceptable-use policy, training
Moderate	Internal summaries, knowledge search, analysis	Use-case approval, data limits, logging, sampled review
High	Customer-facing content, decisions about people, regulated processes	Full assessment, bias and explainability testing, human review, vendor review, monitoring, annual re-approval

Tier	Examples	Minimum controls
Critical	Systems that act without per-action approval	High-tier controls, least-privilege identity, action limits, real-time monitoring, tested stop procedure
Prohibited	Uses outside risk appetite or law	Not approved

FIVE QUESTIONS LEADERSHIP SHOULD BE ABLE TO ANSWER

1. Which material AI systems operate in or interact with the organization?
2. Who owns each system's purpose, controls, technical environment and re-review?
3. What can each system access, influence and do today?
4. What evidence shows each system stays within its approved role?
5. What would cause us to reassess, restrict, pause or retire it?

Common Mistakes

- Governing at launch and then stopping.
- Handing the whole job to a vendor or consultant; accountability stays with the organization.
- Banning AI instead of governing it. Unapproved use continues without visibility.
- Calling human review a control without testing whether it is substantive.
- Governing the model instead of the full system: data, tools, permissions, authority and monitoring.