AI Governance Framework

A practical, board■ready governance model for enterprise AI.

Version	2026 Edition	**Date**	January 07, 2026

This document is provided for informational purposes and does not constitute legal advice. Adapt the framework to your organization, risk profile, and regulatory environment.

If you'd like help implementing this framework, contact: advisory@cichocki.com

Confidential — For internal executive use

Executive summary

AI can create outsized value—while introducing new classes of operational, legal, security, and reputational risk. A governance framework ensures that AI initiatives move fast without increasing unmanaged exposure. This framework is designed to be lightweight, implementable, and defensible in executive and board settings.

Core principles

- Accountability: assign clear ownership for approvals, outcomes, and incidents.
- Transparency: document key decisions, data sources, and limitations in plain language.
- Security & privacy: protect data, control access, and continuously monitor for leakage or abuse.
- Risk■based controls: apply stricter review and monitoring as impact increases (tiering).
- Value discipline: fund initiatives with measurable outcomes and stop low■ROI efforts early.

Governance operating model

A clear operating model separates strategic oversight (board/executives) from day■to■day controls (council/teams). Use this as a starting point and adapt titles to match your organization.

Role	Responsibilities
Board / Audit Committee	Oversight, risk appetite, accountability; receives quarterly AI risk and value reporting.
Executive Sponsor	Sets priorities, resolves conflicts, ensures funding and cross■functional alignment.
AI Governance Council	Approves high■impact use cases, policies, tiering rules; tracks the portfolio.
Risk/Compliance/Legal	Defines controls, reviews high■risk uses, ensures regulatory and contractual compliance.
Product / Engineering	Builds and operates AI systems; maintains documentation, monitoring, and incident response.
Data Governance	Data quality, lineage, and access controls; ensures proper data use and retention.

Decision rights

- Tier 1 (Low impact): team■level approval with standard controls.
- Tier 2 (Medium impact): governance council review, formal documentation, baseline monitoring.
- Tier 3 (High impact): executive approval, enhanced testing, legal/compliance sign■off, and ongoing reporting.

Minimum viable policy pack

o your

Acceptable Use

What AI tools/models are approved, prohibited, and permitted with restrictions.

Data Use & Privacy

What data can be used, retention rules, and sensitive data handling.

Model Risk Management

Validation requirements, bias testing, and documentation expectations.

Vendor & Third■Party

Procurement requirements, security reviews, and contractual protections.

Human Oversight

Where humans must remain in the loop and escalation paths.

Incident Response

Detection, reporting, and remediation for AI■related incidents.

Quick■start checklist

- Define tiering criteria (impact, autonomy, data sensitivity, user exposure).
- Establish required artifacts per tier (model card, data sheet, test plan).
- Create an approval workflow with clear decision rights and SLAs.
- Implement monitoring: drift, hallucination rates, incidents, and value KPIs.

Control gates across the AI lifecycle

Governance works when embedded into delivery. These gates define where controls apply, what evidence [illegible] wh

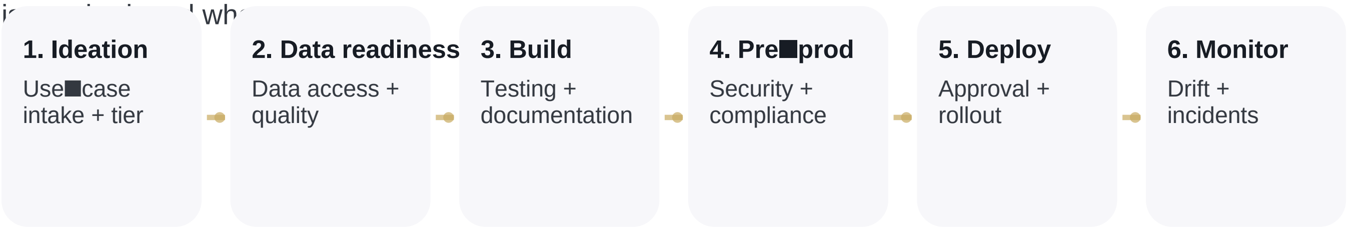

1. Ideation
Use■case intake + tier

2. Data readiness
Data access + quality

3. Build
Testing + documentation

4. Pre■prod
Security + compliance

5. Deploy
Approval + rollout

6. Monitor
Drift + incidents

Evidence artifacts (examples)

- Model card: purpose, training data summary, limitations, intended users, and risks.
- Data sheet: sources, lineage, quality checks, retention, and access controls.
- Test plan: accuracy, bias, robustness, security, and red■team results.
- Approval record: sign■offs, tier, and required mitigations.
- Monitoring plan: metrics, alert thresholds, and incident runbooks.

Implementation roadmap (30/60/90 days)

First 30 days

- Name an executive sponsor and form a small governance council.
- Define tiering criteria and minimum artifacts.
- Publish a one■page acceptable■use and vendor guardrails policy.
- Create a single intake form and approval workflow.

Days 31–60

- Implement monitoring for priority systems (drift, incidents, value KPIs).
- Establish incident response playbook and escalation paths.
- Integrate controls into delivery pipelines (gates + checklists).
- Begin quarterly reporting to executives and the board.

Days 61–90

- Expand the policy pack and training program for teams.
- Harden vendor due diligence and contractual controls.
- Introduce periodic audits for high■impact systems.
- Iterate governance based on metrics and incidents.