

AI Use Case Governance Register

AI Amplification Layer | Reusable template for governing AI use cases from idea through retirement

Use this AI Use Case Governance Register to make every AI initiative visible, owned, classified, controlled, monitored, and reviewable. The register prevents AI work from becoming scattered pilots, hidden automation, vendor-led experimentation, or unmanaged agent behavior. It turns each AI use case into a governed operating-model object tied to ownership, data boundaries, decision rights, human review, controls, risk acceptance, evidence, monitoring, value, and lifecycle state.

1. Object metadata

Metadata	Value
Object type	LPM Knowledge Object
Primary LPM layer	AI Amplification
Connected layers	Ownership Map, Decision Architecture, Information Ecology, Platform Structure, Governance Architecture, Control Map, Risk Acceptance Register, AI Governance Checklist, Agent Accountability Checklist, Human-in-the-Loop Model
Primary use	Maintain a governed register of AI use cases from idea through retirement, with ownership, risk, permissions, evidence, controls, monitoring, and value realization.
Website use	Downloadable register, AI governance artifact, design partner worksheet, JSON object for Lapemo ingestion, and future guided skill.
Version	1.0
Owner	LPM / Lapemo
Last reviewed	2026-06-24

2. Core principles

Principle	Meaning
Register before scale	Every AI use case should be visible before it becomes an operational dependency.
One owner, many contributors	The register may include technical, data, risk, and control owners, but one business owner is accountable for the use case outcome.
Govern the strongest action	Classify the use case by the most powerful thing AI can do, not the marketing description or average workflow.
Data access is a decision right	AI access to sources, APIs, tools, and systems is part of the operating model and must be approved.
Human review must match impact	Human-in-the-loop rules must be based on risk, autonomy, sensitivity, and reversibility.
Evidence beats enthusiasm	The register must connect use cases to evidence, validation, monitoring, and value signals.
Lifecycle governance is mandatory	Use cases should move through states, reviews, restrictions, supersession, and retirement.
AI cannot be invisible work	AI pilots, agents, prompts, vendors, workflows, and automations need traceable ownership and governance.

3. Required register fields

Field	Definition	Required
Use case ID	Unique identifier for the AI use case	Yes
Use case name	Plain-language name of the AI capability	Yes
Business purpose	Outcome, workflow, decision, risk, customer, employee, or control need supported	Yes
Accountable business owner	Person or role accountable for outcome, value, and risk acceptance	Yes
Technical owner	Person or role accountable for build, integration, reliability, and logs	Yes
Data or knowledge owner	Person or role accountable for sources, data quality, sensitivity, access, and freshness	Yes
Control owner	Person or role accountable for control design, evidence, testing, monitoring, and exceptions	Required for moderate and above
AI type	Assistant, copilot, summarizer, classifier, recommender, RAG workflow, automation, model, agent, or decision support	Yes
Workflow or decision supported	Named workflow, decision, product, service, function, or control connected to the AI use case	Yes
Impact tier	Low, moderate, high, critical, regulated, customer-facing, employee-impacting, financial, security, privacy, or control-impacting	Yes
Autonomy level	Draft only, recommend, route, trigger, update, execute with approval, execute within guardrails, or autonomous action	Yes
Approved sources	Systems, documents, dashboards, data products, APIs, and knowledge objects AI may use	Yes
Blocked sources	Sensitive, stale, unapproved, private, confidential, draft, or unsupported sources AI may not use	Yes
Allowed actions	What AI may retrieve, draft, summarize, score, route, trigger, update, or execute	Yes
Blocked actions	What AI may not do without approval or stronger governance	Yes
Human-in-the-loop rule	Reviewer, trigger, authority, evidence package, override, and escalation path	Yes
Controls and evidence	Controls, logs, tests, validations, approvals, screenshots, monitoring, and audit artifacts	Yes
Risk acceptance link	Risk acceptance ID or statement that no acceptance is required	Required for moderate and above
Monitoring signals	Quality, drift, incidents, usage, adoption, value, rejection, override, and control-failure signals	Yes
Lifecycle status	Idea, intake, discovery, design, pilot, limited release, production, scaled, restricted, paused, retired, or superseded	Yes
Review cadence	Weekly, monthly, quarterly, release-based, incident-based, source-change, model-change, or policy-change based	Yes

4. Lifecycle states

Status	Meaning	Governance requirement
Idea	Potential AI opportunity is visible but not yet assessed	Log owner, purpose, expected workflow, and initial risk guess
Intake	Use case submitted for evaluation	Validate business owner, AI type, data boundary, and expected value

Status	Meaning	Governance requirement
Discovery	Feasibility and risk are being assessed	Map sources, decision rights, control needs, and human review rules
Design	Operating model and technical path are being defined	Define permissions, controls, monitoring, evidence, escalation, and launch criteria
Pilot	Limited test with restricted users, sources, and actions	Track outcomes, incidents, failure modes, review burden, and evidence quality
Limited release	Controlled production use with narrow scope	Confirm controls, owner review, monitoring, data boundaries, and support model
Production	Operational use case becomes part of real work	Run recurring governance review and maintain logs, value, risk, and control evidence
Scaled	Use case expands across functions, regions, workflows, or systems	Reassess ownership, decision rights, integration risk, and governance capacity
Restricted or paused	Use case is limited due to risk, failure, stale data, incidents, or control weakness	Document reason, owner, remediation, risk acceptance, and restart criteria
Retired or superseded	Use case is removed or replaced	Record successor, archival evidence, deleted access, and communication plan

5. Classification model

Dimension	Question	Examples
Impact tier	What could be affected?	Customer, employee, financial, operational, compliance, security, privacy, reputation, control, or executive decision impact
Autonomy level	What can AI do?	Retrieve, summarize, draft, recommend, score, route, trigger, update, execute, communicate, or act autonomously
Data sensitivity	What does AI use?	Public, internal, confidential, regulated, personal, financial, HR, customer, security, source code, legal, or privileged data
Decision influence	Does AI shape a decision?	No decision, low-risk recommendation, material decision support, approval workflow, regulated decision, or executive decision
System access	Can AI touch systems?	Read only, write with approval, write within guardrails, tool calls, API updates, production changes, or external communication
Reversibility	Can harm be undone?	Easy to reverse, manual correction needed, customer/employee visible, financial exposure, legal exposure, or irreversible impact
Control dependency	What controls apply?	Access control, evidence review, privacy, security, model validation, audit log, change control, incident management, or kill switch

6. Company-scale versions

Company scale	Register design	Ownership pattern	Governance expectation
500 employees	Single enterprise register or one register per major function. Keep fields lightweight but mandatory for owner, purpose, sources, allowed actions, human review, and status.	Named executive sponsor, accountable business owner, technical owner, data/knowledge owner, and control owner for higher-risk use cases.	Weekly or biweekly review for active pilots. Governance can be centralized, but side-channel pilots should be blocked.
5,000 employees	Federated register by function, product group, region, or platform, with one enterprise rollout. Standardize IDs, statuses, tiers, owners, and evidence links.	Business owner plus domain AI steward, platform owner, data owner, risk/control owner, and review forum for higher-impact use cases.	Monthly portfolio governance, escalation for cross-functional use, and required risk acceptance for high-impact or exception-based use.
10,000+ employees	Enterprise AI use case registry with federated sub-registers, automated ingestion from intake,	Named business owner, accountable executive, domain steward, model/platform owner,	Formal lifecycle gates, audit trails, policy-based classification, control evidence, exception

Company scale	Register design	Ownership pattern	Governance expectation
	platform, vendor, agent, workflow, model, and risk systems.	data owner, control owner, privacy/security/legal/risk owner when required.	management, automated monitoring, and board/executive visibility for critical use.

7. Starter register template

Use case ID	Name	AI action level	Impact tier	Approved sources	Blocked action	Accountable owner	Human review rule	Lifecycle status
AI-001	Sales proposal copilot	Draft only	Moderate	CRM, pricing guide, approved case studies	Cannot send externally without approval	Sales Ops Director	Human approval before customer send	Pilot
AI-002	Policy Q&A assistant	Retrieve and summarize	Moderate	Approved policy knowledge base	Cannot cite stale or draft policy	HR Operations Owner	Escalate conflicts to policy owner	Limited release
AI-003	Invoice exception triage	Classify and route	High	ERP, invoice images, vendor master	Cannot approve payment	Finance Operations Owner	Control owner review for thresholds	Discovery
AI-004	Customer support response agent	Draft and recommend	High	Support KB, CRM case history	Cannot promise refunds or contract exceptions	Service Owner	Approval for sensitive responses	Design
AI-005	Executive risk summary generator	Summarize	High	Risk register, control evidence, incident logs	Cannot create new risk acceptance	Risk Governance Owner	Evidence checklist required	Production

8. Governance readiness scoring

Score each use case from 0 to 100 before pilot, launch, scale, and renewal. High-impact use cases should not proceed with weak ownership, undefined source boundaries, missing control evidence, or unclear human review.

Score area	Points	What good looks like
Ownership clarity	0-20	Named business, technical, data/knowledge, decision, and control owners
Use case classification	0-15	AI type, impact tier, autonomy, data sensitivity, system access, and reversibility are classified
Data and source boundary	0-15	Approved and blocked sources are defined with source-of-truth and freshness rules
Decision and action boundary	0-15	Allowed and blocked AI actions are explicit and tied to decision rights
Human review and escalation	0-15	Human-in-the-loop trigger, reviewer authority, override, stop, and escalation rules are documented
Controls and evidence	0-10	Controls, logs, validation, monitoring, and evidence package are sufficient for impact tier
Value and lifecycle management	0-10	Value metric, review cadence, lifecycle state, retirement criteria, and supersession path are present

9. Validation rules

Rule	Requirement
Owner required	Block use case from pilot if no accountable business owner is named.
Impact tier required	Block governance approval if impact tier, autonomy level, data sensitivity, or system access are blank.
Moderate-plus control rule	Moderate, high, critical, regulated, customer-facing, employee-impacting, financial, security, privacy, and control-impacting use cases require control ownership.
Source boundary rule	Use case cannot launch if approved sources and blocked sources are not listed.
Write-action rule	Any AI use case with write access, external communication, payment, entitlement, workflow execution, or system update needs explicit approval and control evidence.
Human review rule	High-impact or irreversible outcomes require named reviewer, review trigger, evidence package, stop rule, and escalation path.
Risk acceptance rule	Accepted residual risk must have owner, expiration date, evidence, compensating controls, and review cadence.
Monitoring rule	Production and scaled use cases require performance, quality, drift, incident, override, adoption, and value monitoring.
Supersession rule	Retired, paused, restricted, or superseded use cases must preserve decision and evidence trail.

10. Mapping rules

LPM object or system	Mapped fields
Ownership Map	Accountable business owner, technical owner, data owner, control owner, executive sponsor
Decision Rights Model	AI action boundary, decision owner, approval authority, escalation path
Human-in-the-Loop Model	Review trigger, reviewer authority, evidence package, override and stop rules
Source-of-Truth Map	Approved sources, blocked sources, system of record, freshness, lineage
Data Lineage Map	Data movement, transformation, model input, output, dashboard, decision, and control evidence
Control Map	Preventive, detective, corrective, access, privacy, security, and operational controls
Risk Acceptance Register	Accepted residual risk, expiration, compensating controls, and review obligations
Agent Accountability Checklist	Agent owner, permissions, tools, memory, logs, incidents, and lifecycle
Lapemo platform	Register row becomes an ingestible control-plane object with status, score, owner, risk, evidence, and lifecycle links

11. AI prompt starters

Prompt purpose	Prompt
Classify a use case	Given this AI use case description, classify AI type, impact tier, autonomy level, data sensitivity, system access, reversibility, and required governance path.
Find missing fields	Review this AI use case register row and identify missing ownership, decision rights, source, control, evidence, monitoring, and lifecycle fields.
Generate a launch gate	Create a launch gate checklist for this AI use case based on its impact tier, allowed actions, data sources, and human-in-the-loop rule.

Prompt purpose	Prompt
Draft risk questions	Generate governance questions that risk, legal, security, privacy, architecture, and control owners should answer before this use case proceeds.
Recommend lifecycle status	Based on this use case evidence, incidents, monitoring signals, and value metrics, recommend whether it should remain in pilot, move to production, pause, restrict, scale, retire, or supersede.

12. Operating cadence and reuse model

- Intake review: weekly for active design partners or AI-heavy functions.
- Governance review: monthly for active pilot and production use cases.
- Executive review: quarterly for scaled, high-impact, regulated, or cross-enterprise AI use.
- Triggered review: required after model change, vendor change, source change, policy change, incident, control failure, significant drift, or scope expansion.
- Reuse model: this artifact should exist as a human guide, downloadable template, machine-readable JSON object, and guided Lapemo skill. Systems can recommend updates, but accountable humans approve changes.