AI Governance Checklist

Reusable LPM Knowledge Object for governance-layer AI ownership, evidence, controls, data boundaries, and lifecycle review

Use this AI Governance Checklist to decide whether an AI use case is safe, owned, evidenced, controlled, monitored, and ready to proceed. The checklist prevents AI work from moving forward as disconnected pilots, shadow automation, vendor enthusiasm, or executive theater. It ties each AI initiative to accountable ownership, decision rights, data boundaries, evidence standards, control coverage, human review, monitoring, and lifecycle governance.

Built for three company profiles: 500+ employees, 5,000+ employees, and 10,000+ employees. Use it as a website artifact, AI governance review guide, pilot readiness checklist, and future Lapemo guided skill.

Metadata

Metadata	Value
Object type	LPM Knowledge Object
Primary LPM layer	Governance Architecture
Connected layers	Ownership Map, Decision Architecture, Information Ecology, Platform Structure, Control Map, Risk Acceptance Register, AI Amplification
Primary use	Validate whether an AI initiative, model, agent, workflow, or AI-assisted operating process is governed enough to design, pilot, launch, scale, or retire.
Website use	Downloadable governance checklist, AI readiness resource, executive governance guide, JSON object for Lapemo ingestion, and future guided skill.
Version	1.0
Owner	LPM / Lapemo
Last reviewed	2026-06-24

Core principles

Principle	Meaning
AI must have an accountable owner	Every AI use case, model, agent, workflow, prompt, or automation must have one accountable business owner and clear supporting technical, data, risk, and control owners.
AI cannot outrun decision rights	AI may recommend, draft, retrieve, summarize, score, route, trigger, or act only within explicit decision rights and approval boundaries.
Data boundaries are governance boundaries	AI must only use approved sources, approved data classes, defined retrieval scope, and clear source-of-truth rules.
Evidence must beat confidence	AI confidence, fluent summaries, or model scores are not enough. Claims must connect to source, freshness, lineage, owner, and validation evidence.
Human review must be designed, not assumed	Human-in-the-loop rules must specify who reviews, when review happens, what evidence is reviewed, and what authority the reviewer has.
Controls must be mapped before scale	Controls, monitoring, logs, thresholds, overrides, escalation paths, and kill-switch ownership must be defined before AI impacts customers, employees, financials, compliance, or enterprise decisions.
AI governance must be lifecycle-based	Governance does not stop at launch. AI use must be reviewed, monitored, updated, superseded, restricted, or retired as data, models, vendors, policies, and operating conditions change.

Required fields

Field	Definition	Required
AI use case ID	Unique identifier for the AI initiative, model, agent, automation, workflow, prompt library, or AI-assisted process	Yes
AI use case name	Plain-language name of the AI use case or governed AI capability	Yes
Business purpose	Outcome, problem, decision, workflow, control, or customer/employee need the AI use case supports	Yes
Accountable business owner	Person or role accountable for outcomes, risk acceptance, value realization, and ongoing use	Yes
Technical owner	Person or role accountable for implementation, integration, reliability, logs, and operational health	Yes
Data owner	Person or role accountable for approved sources, data quality, sensitivity, access, lineage, and freshness	Yes
Control owner	Person or role accountable for controls, evidence, monitoring, testing, and exception handling	Yes
Decision owner	Person or forum with authority over AI-enabled decisions or recommendations	Yes
AI type	Assistant, copilot, classifier, summarizer, recommender, retrieval workflow, automation, model, agent, decision support, or autonomous action	Yes
Impact tier	Low, moderate, high, critical, regulated, customer-facing, employee-impacting, financial, security, privacy, or control-impacting	Yes
Approved data sources	Systems, documents, dashboards, data products, knowledge objects, repositories, and APIs the AI may use	Yes
Blocked data sources	Sources, sensitive data, unapproved systems, stale documents, private channels, or unsupported knowledge bases the AI may not use	Yes
Allowed AI actions	What AI is allowed to retrieve, draft, recommend, score, route, trigger, update, communicate, monitor, or execute	Yes
Blocked AI actions	What AI may not do without human approval, governance approval, or control validation	Yes
Human review rule	Reviewer, review trigger, review evidence, decision authority, override path, and exception handling	Yes
Evidence requirement	Sources, logs, tests, model cards, vendor documentation, policy references, validation results, and decision records required	Yes
Control coverage	Preventive, detective, corrective, access, privacy, security, compliance, model, vendor, and operational controls	Yes
Monitoring signals	Performance, drift, incidents, usage, adoption, false positives, false negatives, control failures, user feedback, and business value metrics	Yes
Escalation path	Where issues go when AI output is wrong, harmful, unauthorized, stale, risky, overused, or outside approved scope	Yes
Risk acceptance link	Accepted risk ID or explicit statement that no risk acceptance is required	Required for moderate and above
Launch status	Idea, discovery, design, pilot, limited release, production, scaled, restricted, paused, retired, or superseded	Yes
Review cadence	Weekly, monthly, quarterly, release-based, incident-based, policy-change based, vendor-change based, or data-change based	Yes

Governance domains

Domain	Purpose
Business ownership	Confirms the AI use case has a real business owner, value case, accountable outcome, and funding path.
Decision rights	Defines whether AI advises, drafts, routes, recommends, scores, approves, updates, or acts, and who has authority over each action.
Data and knowledge boundaries	Defines approved sources, blocked sources, source-of-truth rules, sensitivity, access, freshness, and lineage.
Model, vendor, and tool governance	Documents the AI tool, model, vendor, platform, version, terms, data handling, reliability, and lifecycle ownership.
Human review and accountability	Defines required review, approval, override, exception, and escalation paths for AI-assisted work.
Controls and evidence	Maps controls, logs, test evidence, evidence owners, auditability, monitoring, and control failure handling.
Risk, legal, security, and privacy	Confirms required reviews for sensitive data, customer impact, employee impact, regulated decisions, IP, security, and privacy.
Operational monitoring	Defines performance signals, drift monitoring, incidents, adoption, business value, user feedback, and retirement triggers.

AI governance checklist

1. Business ownership and intent

Checklist item	Status	Guidance
Use case has one accountable business owner	Yes / No / Partial	Do not proceed without a named owner.
Business purpose is tied to a measurable outcome	Yes / No / Partial	Outcome should be specific enough to measure value or harm.
AI use case is connected to a workflow, decision, service, product, control, or operating-model need	Yes / No / Partial	Avoid standalone AI demos with no operating destination.
Funding, support, and lifecycle ownership are clear	Yes / No / Partial	Owner must cover sustainment, not just launch.

2. Decision rights and action boundary

Checklist item	Status	Guidance
AI role is classified as retrieve, draft, summarize, recommend, score, route, trigger, update, or act	Yes / No / Partial	Classify the strongest action AI can take.
Decision owner is documented for every AI-influenced decision	Yes / No / Partial	AI cannot become the hidden decision owner.
Allowed AI actions are explicitly defined	Yes / No / Partial	Allowed actions must be narrower than tool capability.
Blocked AI actions are explicitly defined	Yes / No / Partial	Block approvals, commitments, external communications, financial actions, policy exceptions, or control bypass unless approved.
Escalation path exists when AI output conflicts with human judgment, policy, evidence, or controls	Yes / No / Partial	Conflicts should not be solved in side channels.

3. Data, source, and knowledge boundary

Checklist item	Status	Guidance
Approved sources are named and owned	Yes / No / Partial	List systems, documents, dashboards, repositories, APIs, and knowledge objects.
Blocked sources and sensitive data classes are named	Yes / No / Partial	Include personal, confidential, regulated, stale, draft, private, and unsupported sources.
Source-of-truth rule is documented	Yes / No / Partial	Define what source wins when sources conflict.
Freshness window is defined for AI-readable knowledge	Yes / No / Partial	Stale knowledge should trigger warning, review, or blocked use.
Data lineage and evidence traceability are sufficient for the use case impact tier	Yes / No / Partial	Higher impact use cases need stronger lineage.

4. Model, vendor, and platform governance

Checklist item	Status	Guidance
Model, vendor, tool, platform, and version are documented	Yes / No / Partial	Do not govern AI generically when the implementation has specific behavior.
Data handling, retention, training use, logging, and access terms are understood	Yes / No / Partial	Confirm vendor and enterprise terms.
Security, privacy, legal, procurement, and architecture review needs are classified	Yes / No / Partial	Do not assume all AI tools fit the same path.
Model limitations, known failure modes, and misuse risks are documented	Yes / No / Partial	Use plain language that owners can understand.

5. Human review and operating control

Checklist item	Status	Guidance
Human review rule is specific	Yes / No / Partial	Name reviewer role, trigger, evidence, authority, and override path.

Control owner is documented	Yes / No / Partial	Control ownership cannot be split across everyone.
Preventive, detective, and corrective controls are mapped	Yes / No / Partial	At minimum define guardrails, logs, monitoring, and remediation.
Kill switch, pause rule, or rollback path exists	Yes / No / Partial	Critical AI use needs a clear stop path.

6. Risk, compliance, and ethics review

Checklist item	Status	Guidance
Impact tier is assigned	Yes / No / Partial	Higher impact requires stronger review and evidence.
Customer, employee, financial, compliance, security, privacy, or regulatory impact is classified	Yes / No / Partial	Material impact changes the governance route.
Risk acceptance is documented if residual risk remains	Yes / No / Partial	Accepted risk must be explicit, owned, time-bound, and reviewed.
Bias, harm, misuse, overreliance, and explainability risks are considered	Yes / No / Partial	Use practical risk language, not theory only.

7. Evidence, logs, and monitoring

Checklist item	Status	Guidance
Evidence package is complete enough for impact tier	Yes / No / Partial	Evidence should include source, tests, logs, approvals, validation, and operating assumptions.
Output validation method is documented	Yes / No / Partial	Define how wrong, incomplete, stale, or harmful outputs are detected.
Monitoring signals are defined	Yes / No / Partial	Track accuracy, drift, incidents, overrides, user feedback, adoption, value, and control failures.
Decision logs or audit records are produced where required	Yes / No / Partial	Critical AI decisions need replayability.

8. Launch, scale, and lifecycle governance

Checklist item	Status	Guidance
Launch gate is defined	Yes / No / Partial	Idea, discovery, pilot, production, scaled, restricted, paused, or retired.
Scale criteria are defined before broad rollout	Yes / No / Partial	Pilot success does not automatically mean enterprise scale.
Review cadence and trigger conditions are documented	Yes / No / Partial	Review on incidents, data change, policy change, vendor change, model change, drift, and value degradation.
Retirement or supersession rule is documented	Yes / No / Partial	AI capabilities must be removable when unsafe, stale, redundant, or low value.

Version for 500+ employee company

Dimension	Recommended pattern
Design intent	Create a practical AI governance checkpoint that prevents unmanaged pilots while keeping review lightweight and fast.
Minimum scope	Track business owner, use case, approved sources, AI actions, human review, risks, controls, evidence, and launch status.
Operating pattern	A small cross-functional AI review group reviews moderate and high-impact use cases before pilot or production use.
AI focus	Focus on ownership, approved tool use, source boundaries, human review, and preventing shadow automation.
Governance need	Connect to ownership map, decision rights model, source-of-truth map, control map, and risk acceptance register.
Red flags	Teams use public AI tools without data rules, pilots have no owner, and leaders measure activity instead of business outcome.

Version for 5,000+ employee company

Dimension	Recommended pattern
Design intent	Create federated AI governance where business units can move quickly inside enterprise policy, controls, evidence, and review rules.
Minimum scope	Add impact tiering, data classification, model/tool inventory, vendor review, logs, monitoring, escalation, and risk acceptance linkage.
Operating pattern	Business unit AI owners submit use cases through a common checklist, with routing to security, privacy, legal, data, risk, architecture, and governance as needed.
AI focus	Focus on preventing duplicate pilots, conflicting standards, unsafe integrations, weak evidence, and unsupported production AI.
Governance need	Connect to platform map, integration map, data lineage map, evidence checklist, decision log, and control map.
Red flags	Different functions create different AI rules, review is slow and unclear, and AI tools are approved without operational monitoring.

Version for 10,000+ employee company

Dimension	Recommended pattern
Design intent	Create enterprise AI governance that can scale across business units, regions, vendors, regulated functions, shared platforms, and autonomous agents.
Minimum scope	Full lineage across use case, owner, model/tool, data, source, control, risk, decision, approval, logs, incidents, value, and retirement path.
Operating pattern	Central AI governance defines policy, risk tiering, standards, and control requirements while federated owners manage intake, evidence, monitoring, and lifecycle review.
AI focus	Focus on high-impact systems, external communication, employee or customer decisions, regulated workflows, agent action, sensitive data, and control bypass risk.
Governance need	Integrate with GRC, model risk, data governance, privacy, security, architecture, procurement, internal audit, enterprise risk, and executive governance.
Red flags	AI is embedded in workflows without traceability, vendors change models without review, agents act across systems without ownership, and audit cannot replay decisions.

Scoring logic

Dimension	Score	What good looks like
Business ownership	0-5	A real owner is accountable for purpose, outcome, risk, value, lifecycle, and escalation.
Decision boundary clarity	0-5	Allowed and blocked AI actions are explicit and tied to decision rights.
Data boundary strength	0-5	Approved sources, blocked sources, source-of-truth rules, sensitivity, lineage, and freshness are defined.
Model and vendor governance	0-5	Tool, model, version, vendor terms, limitations, failure modes, and lifecycle ownership are documented.
Human review design	0-5	Review triggers, reviewer role, evidence, approval authority, override path, and exception handling are clear.
Control coverage	0-5	Preventive, detective, corrective, access, privacy, security, operational, and AI controls are mapped.
Evidence strength	0-5	Evidence is current, sourced, owned, confidence-rated, testable, and sufficient for the impact tier.
Risk classification	0-5	Impact tier, residual risk, compliance obligations, risk acceptance, and escalation are clear.

Monitoring readiness	0-5	Performance, drift, incidents, usage, value, user feedback, and control failures are monitored.
Lifecycle discipline	0-5	Launch, scale, review, restriction, pause, retirement, and supersession rules are documented.

Suggested readiness score: average the ten scores, then classify 0-1.9 as Ungoverned or hidden AI activity, 2.0-3.4 as Documented but weak AI governance, 3.5-4.4 as Governed AI operating model, and 4.5-5.0 as AI-ready governance architecture.

AI prompts

- Classify this AI use case by impact tier, affected LPM layers, likely owners, required evidence, governance route, and allowed versus blocked AI actions.
- Review this AI governance checklist and identify missing ownership, weak evidence, unclear data boundary, missing controls, unresolved risk, and undefined human review.
- Given the AI use case description, generate the minimum governance evidence package required before pilot, production launch, or enterprise scale.
- Score this AI use case from 0 to 5 across business ownership, decision boundary clarity, data boundary strength, model and vendor governance, human review design, control coverage, evidence strength, risk classification, monitoring readiness, and lifecycle discipline.
- Determine whether this AI use case should proceed locally, require domain review, require enterprise governance, require executive approval, require risk acceptance, or be blocked.
- Convert this AI governance checklist into Lapemo objects for owner, use case, decision, data source, control, evidence, platform, integration, risk, monitoring signal, and AI boundary.

Validation rules

- Every AI use case must have one accountable business owner before pilot or launch.
- Every AI use case must classify AI role and strongest allowed action.
- Every AI use case must define allowed AI actions and blocked AI actions.
- Every AI use case must define approved data sources and blocked data sources.
- Every moderate, high, critical, regulated, customer-facing, employee-impacting, financial, security, privacy, or control-impacting use case must have documented human review.
- Every AI use case must have evidence requirements appropriate to its impact tier.
- Every AI use case must identify control owner, monitoring signals, escalation path, and stop or rollback path where material risk exists.
- Any residual risk must be linked to a risk acceptance register entry or explicitly marked as no acceptance required.
- AI cannot approve, commit, externalize, update systems of record, bypass controls, or make regulated decisions unless those actions are explicitly approved by governance.
- Production AI use must have review cadence, trigger conditions, lifecycle status, and retirement or supersession rule.
- Status must never be encoded only by color. Use labels such as idea, discovery, pilot, production, scaled, restricted, paused, retired, or superseded.

Lapemo ingestion mapping

Lapemo object	Fields / entities	Use
Knowledge object	AI Governance Checklist	Canonical reusable artifact for governance-layer AI review and operating control.
AI use case object	Use case ID, name, purpose, AI type, impact tier, lifecycle status	Creates structured AI initiative records.
Ownership object	Business owner, technical owner, data owner, control owner, decision owner, escalation owner	Connects AI to accountable owners.
Decision object	Decision rights, allowed actions, blocked actions, approval authority, launch decision, scale decision	Connects AI use to decision architecture.

Data object	Approved sources, blocked sources, sensitivity, lineage, freshness, source-of-truth rule	Controls what AI can retrieve or use.
Control object	Preventive controls, detective controls, corrective controls, monitoring, failure handling, kill switch	Connects AI activity to governance controls.
Evidence object	Evidence package, tests, logs, source, owner, freshness, confidence, approval record	Creates audit-ready proof.
Platform object	AI tool, model, vendor, platform, integration, system of record, workflow	Connects AI governance to platform structure.
Risk object	Impact tier, residual risk, risk acceptance link, compliance impact, privacy impact, security impact	Routes risk to the correct governance path.
Monitoring object	Performance, drift, incidents, usage, adoption, business value, control failures, review triggers	Tracks whether AI remains safe and valuable after launch.

Website positioning

Use this artifact as a downloadable governance-layer AI governance checklist and as a future guided skill inside Lapemo. The website version should explain that AI governance is not policy theater. It is operating-model control for ownership, data, decisions, evidence, controls, risk, monitoring, and lifecycle management.

Reusable knowledge object structure

Layer	Reusable asset
Human guide	Plain-English explanation, governance domains, checklist sections, scale versions, and scoring logic.
Downloadable template	DOCX and PDF for governance reviews, AI readiness workshops, pilot assessments, and executive briefings.
Machine-readable schema	JSON object with required fields, scoring logic, prompts, validation rules, and Lapemo mappings.
Guided skill	Future Lapemo workflow that asks governance questions, scores readiness, validates evidence, checks controls, and recommends proceed, review, escalate, block, or retire.