Risk Acceptance Register

Reusable LPM Knowledge Object for governance-layer risk acceptance, evidence, controls, ownership, and AI boundaries

Use this Risk Acceptance Register to make accepted risk explicit, owned, evidence-backed, time-bound, and reviewable. The register prevents risk acceptance from becoming a vague meeting note, email approval, informal exception, or permanent workaround. It is designed for companies scaling AI where accepted risk must be tied to business outcomes, decision rights, controls, evidence, systems, data, AI boundaries, and review triggers.

Built for three company profiles: 500+ employees, 5,000+ employees, and 10,000+ employees. Use it as a website artifact, governance workshop guide, risk acceptance template, 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, AI Amplification
Primary use	Document accepted risks with accountable ownership, rationale, evidence, controls, expiration dates, review triggers, and escalation paths.
Website use	Downloadable governance template, AI readiness resource, risk workshop 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
Accepted risk must be explicit	A risk is not accepted because nobody fixed it. It is accepted only when a defined authority documents the rationale, owner, evidence, expiration, and review rule.
Risk acceptance is not risk deletion	Accepted risk remains visible until it is mitigated, transferred, retired, escalated, superseded, or formally renewed.
Ownership must be singular	Each accepted risk needs one accountable business or executive owner, even when several teams support mitigation.
Evidence must support the decision	The acceptance rationale must include evidence, alternatives considered, residual risk, control coverage, and known tradeoffs.
Time limits protect the enterprise	Every accepted risk needs an expiration date, renewal rule, review cadence, and trigger conditions.
AI raises the bar	AI-related risk acceptance must define allowed actions, blocked actions, human review, audit logs, retrieval limits, model or agent boundaries, and kill-switch ownership.
Exceptions must not become shadow policy	Repeated acceptance of the same risk should trigger control redesign, process redesign, funding decision, or executive governance review.

Required fields

Field	Definition	Required
Risk ID	Unique identifier for the accepted risk, exception, control gap, AI boundary issue, or governance decision	Yes
Risk name	Plain-language name of the risk being accepted	Yes
Risk statement	Specific condition, cause, and potential impact being accepted	Yes
Business context	Why the risk exists now, what objective or delivery need is affected, and why acceptance is being requested	Yes
Impacted objective	Outcome, initiative, customer journey, compliance obligation, AI use case, workflow, system, or operating-model goal affected by the risk	Yes
LPM layer affected	Ownership, Decision Architecture, Communication Architecture, Information Ecology, Platform Structure, Governance Architecture, AI Amplification	Yes
Risk category	Operational, strategic, financial, compliance, regulatory, security, privacy, vendor, data, AI/model, reputational, transformation, or control risk	Yes
Inherent risk rating	Risk level before mitigation or controls	Yes
Residual risk rating	Risk level after existing controls, mitigations, and boundaries	Yes
Acceptance rationale	Why the residual risk is acceptable for a limited period	Yes
Alternatives considered	Mitigate now, defer, transfer, avoid, redesign, automate, escalate, stop work, or accept temporarily	Yes
Accountable risk owner	Person or role accountable for the accepted risk and its outcomes	Yes
Approval authority	Role, forum, committee, executive, or governance body authorized to accept this risk	Yes
Control owner	Owner of the control, mitigation, evidence, monitoring, or compensating control	Required when controls exist
Evidence package	Evidence used to justify acceptance, including source, freshness, confidence, and owner	Yes
Controls and mitigations	Existing controls, temporary controls, compensating controls, monitoring, or operating boundaries	Yes
AI involvement	Whether AI retrieves, recommends, drafts, decides, updates, triggers, monitors, communicates, or creates evidence related to the risk	Yes
AI allowed actions	What AI is permitted to do while the risk is accepted	Required when AI is involved
AI blocked actions	What AI is not permitted to do without human approval or governance review	Required when AI is involved
Human review rule	Human approval, validation, override, escalation, or audit required before action	Required when AI, control, or high risk is involved
Expiration date	Date when acceptance expires and must be closed, renewed, escalated, or redesigned	Yes
Review cadence	Weekly, monthly, quarterly, release-based, incident-based, milestone-based, or governance-cycle review	Yes
Trigger conditions	Events that force earlier review, escalation, closure, renewal, or stop-work	Yes
Escalation path	Where the risk goes if exposure increases, evidence weakens, controls fail, or expiration is missed	Yes
Status	Draft, pending approval, accepted, active monitoring, expired, escalated, remediated, superseded, rejected, or retired	Yes

Risk categories

Category	Purpose	Examples
Operational risk	Risk that work, handoffs, roles, controls, service quality, or execution reliability fail	Unowned workflow, weak handoff, manual control dependency, capacity gap
Decision risk	Risk that authority, evidence, or decision path is unclear or too slow	No decision owner, conflicting approval route, weak evidence, decision debt
Data and information risk	Risk that data, dashboards, knowledge objects, or evidence are inaccurate, stale, incomplete, or uncontrolled	Stale source, unclear data owner, conflicting dashboards, weak lineage
Platform and integration risk	Risk tied to tools, systems, APIs, automations, vendor platforms, or system dependencies	Unsupported integration, duplicate tool, brittle automation, vendor dependency
AI and model risk	Risk that AI outputs, agents, prompts, retrieval, actions, or model behavior create harm or unauthorized decisions	Agent action without owner, weak review, unknown training source, hallucinated evidence
Control and compliance risk	Risk that a control, policy, regulatory obligation, or audit requirement is missing, stale, bypassed, or ineffective	Expired control evidence, policy exception, unresolved audit issue, missing approval
Security and privacy risk	Risk that access, sensitive data, identity, confidentiality, or privacy expectations are breached	Over-permissioned agent, exposed data, weak access control, privacy exception
Financial and reputational risk	Risk that accepted exposure creates financial loss, customer harm, brand damage, or executive accountability exposure	Known customer defect, budget overrun, SLA miss, public trust exposure

Register template

Area	Question	Guidance
Risk identity	What risk is being accepted?	Capture ID, name, risk statement, category, impacted objective, LPM layer, and business context.
Authority model	Who can accept it?	Identify accountable risk owner, approval authority, decision route, governance forum, and executive escalation path.
Evidence model	What proves acceptance is reasonable?	Attach evidence package, source, freshness, confidence, alternatives considered, impact analysis, and control status.
Control model	What keeps the accepted risk bounded?	Document controls, mitigations, compensating controls, monitoring rules, control owner, and failure conditions.
AI boundary	How does AI interact with the risk?	Define allowed AI actions, blocked actions, review rules, logging, retrieval scope, and kill-switch owner.
Lifecycle model	When does acceptance end?	Define expiration date, review cadence, trigger conditions, renewal rule, closure path, and supersession rule.

Workflow

Step	Action
1. Identify risk	Write the risk statement in condition, cause, and impact form. Avoid vague labels such as technical risk or business risk without context.
2. Classify exposure	Assign category, LPM layer, inherent risk, residual risk, impacted objective, affected systems, data, workflows, controls, and AI involvement.
3. Confirm authority	Validate who has decision rights to accept the risk and whether escalation is required by severity, regulation, AI boundary, financial impact, or control gap.
4. Build evidence package	Attach evidence source, freshness, confidence, alternatives considered, control status, mitigation plan, monitoring rule, and tradeoff analysis.
5. Define acceptance terms	Set owner, expiration date, review cadence, trigger conditions, escalation path, compensating controls, and closure or renewal requirements.
6. Approve or reject	Approval authority accepts, rejects, escalates, or requests more evidence. Document the decision in the risk register and decision log.
7. Monitor and review	Track evidence freshness, control performance, incidents, AI boundary violations, expiration dates, and changing business context.
8. Close, renew, or escalate	At expiration or trigger event, close the risk, renew with new evidence, escalate, redesign controls, fund remediation, or stop the affected work.

Version for 500+ employee company

Dimension	Recommended pattern
Design intent	Create simple visibility so accepted risk is not hidden in leadership conversations, spreadsheets, or project notes.
Minimum scope	Track high and material medium risks across critical workflows, customer-impacting work, core systems, security/privacy exposure, and AI initiatives.
Operating pattern	Risk owner, approving executive, basic evidence package, expiration date, monthly review, and clear escalation path.
AI focus	No AI initiative should accept risk without a named business owner, human review rule, data boundary, and stop or rollback path.
Governance need	Use the register with the Control Map, Governance Decision Tree, Decision Log, and AI Initiative Owner Register.
Red flags	Accepted risk has no expiration, no owner, no evidence, no review cadence, or no clear approval authority.

Version for 5,000+ employee company

Dimension	Recommended pattern
Design intent	Standardize risk acceptance across functions, product lines, shared services, vendors, systems, data, and AI use cases.
Minimum scope	Track accepted risks by domain, severity, owner, approval forum, control dependency, evidence source, system dependency, AI involvement, and remediation path.
Operating pattern	Domain-level acceptance with central governance visibility, quarterly executive review, automated expiration alerts, and issue linkage.
AI focus	AI-related accepted risk must include retrieval limits, output review, action permissions, audit logs, model or agent owner, and kill-switch owner.
Governance need	Connect accepted risks to control maps, evidence checklists, platform maps, data lineage maps, integration maps, and decision logs.
Red flags	Different functions use different acceptance rules, exceptions never expire, and AI tools operate under assumed rather than documented risk acceptance.

Version for 10,000+ employee company

Dimension	Recommended pattern
Design intent	Create enterprise-grade risk acceptance lineage across business units, regions, regulated functions, shared services, vendors, AI agents, and critical platforms.
Minimum scope	Track acceptance authority, risk appetite alignment, regulatory exposure, financial exposure, customer impact, policy exceptions, model risk, data risk, and control exceptions.
Operating pattern	Integrated GRC workflow with enterprise risk taxonomy, approvals, evidence packages, expiration monitoring, remediation workflow, and board or executive reporting.
AI focus	AI risk acceptance requires enterprise AI governance review for high-impact systems, automated decisioning, sensitive data, external communication, or material control bypass.
Governance need	Integrate with enterprise risk, compliance, internal audit, legal, privacy, security, model risk, data governance, architecture, finance, procurement, and executive governance.
Red flags	Risk acceptance becomes a workaround for unfunded controls, regional exceptions conflict, accepted AI risk has no monitoring, and audit cannot replay the acceptance rationale.

Scoring logic

Dimension	Score	What good looks like
Risk clarity	0-5	The risk statement clearly states condition, cause, impact, category, LPM layer, and affected objective.
Ownership clarity	0-5	Accountable risk owner, control owner, evidence owner, approval authority, and escalation owner are clear.
Authority fit	0-5	The acceptance decision is made by the correct authority for the risk severity, policy, AI boundary, and business exposure.
Evidence strength	0-5	Evidence is current, sourced, owned, confidence-rated, and sufficient to support the acceptance rationale.
Control coverage	0-5	Controls, mitigations, compensating controls, monitoring, and failure paths are documented and active.
AI boundary clarity	0-5	AI allowed actions, blocked actions, human review, retrieval scope, logs, and kill-switch ownership are explicit.
Expiration discipline	0-5	Acceptance has expiration date, review cadence, trigger conditions, renewal rule, and closure path.
Escalation readiness	0-5	Escalation path is defined for increased exposure, weak evidence, failed controls, incidents, missed expiration, or AI boundary breach.
Remediation path	0-5	There is a clear plan to mitigate, redesign, fund, transfer, retire, or renew the risk.
Auditability	0-5	A reviewer can replay why the risk was accepted, who approved it, what evidence supported it, and what changed after acceptance.

Suggested readiness score: average the ten scores, then classify 0-1.9 as Hidden or informal risk acceptance, 2.0-3.4 as Documented but weakly governed risk acceptance, 3.5-4.4 as Governed risk acceptance register, and 4.5-5.0 as AI-ready risk acceptance architecture.

AI prompts

- Given this risk statement, classify the risk category, affected LPM layer, severity, likely owner, approval authority, evidence requirement, and governance route.
- Review this proposed risk acceptance and identify missing owner, weak evidence, unclear approval authority, missing controls, stale review date, and undefined AI boundary.
- Determine whether this risk should be accepted, mitigated, transferred, avoided, escalated, redesigned, or blocked based on severity, evidence, controls, and AI involvement.
- Score the risk acceptance from 0 to 5 across risk clarity, ownership clarity, authority fit, evidence strength, control coverage, AI boundary clarity, expiration discipline, escalation readiness, remediation path, and auditability.
- Generate the minimum evidence package required before this risk can be accepted by governance.
- Draft acceptance terms including owner, rationale, residual risk, controls, expiration date, review cadence, trigger conditions, escalation route, and closure rule.
- Convert this risk acceptance register row into Lapemo objects for risk, owner, control, evidence, decision, platform, AI boundary, issue, and governance route.

Validation rules

- Every accepted risk must have one accountable risk owner.
- Every accepted risk must have a clear approval authority that matches severity, policy, regulation, financial exposure, customer impact, and AI involvement.
- Every accepted risk must include risk statement, business context, impacted objective, inherent risk, residual risk, and acceptance rationale.
- Every accepted risk must include evidence package, evidence owner, source, freshness, and confidence level.
- Every accepted risk must include expiration date, review cadence, trigger conditions, and closure or renewal rule.
- Every accepted risk must include escalation path and remediation owner.
- Every risk involving AI must define allowed AI actions, blocked AI actions, human review rule, logging rule, retrieval boundary, and kill-switch or stop path.
- Accepted risk cannot remain active past expiration without renewal evidence and renewed approval.
- Repeated acceptance of the same risk must trigger control redesign, funding decision, process redesign, or executive escalation.
- Status must never be encoded only by color. Use labels such as draft, pending approval, accepted, active monitoring, expired, escalated, remediated, superseded, rejected, or retired.

Lapemo ingestion mapping

Lapemo object	Fields / entities	Use
Knowledge object	Risk Acceptance Register	Canonical reusable artifact for governance-layer risk acceptance and monitoring.
Risk object	Risk ID, name, statement, category, inherent risk, residual risk, impacted objective, LPM layer, status	Creates structured risk records that can be monitored and reviewed.
Ownership object	Accountable risk owner, approval authority, control owner, evidence owner, remediation owner, escalation owner	Connects accepted risk to accountable humans and governance forums.
Decision object	Acceptance rationale, alternatives considered, approval decision, decision date, approving forum	Connects risk acceptance to decision logs and decision rights.
Control object	Controls, mitigations, compensating controls, monitoring rules, failure conditions	Links accepted risk to bounded operating controls.
Evidence object	Evidence package, source, owner, freshness, confidence, storage location	Turns acceptance into inspectable proof.
Platform object	Affected systems, integrations, workflows, dashboards, logs, source systems	Connects accepted risk to platforms and system dependencies.
Data object	Data source, sensitivity, lineage, freshness, owner, AI retrieval boundary	Connects accepted risk to information ecology and source-of-truth rules.

AI boundary object	AI involvement, allowed actions, blocked actions, human review, logs, kill switch	Controls AI behavior while risk is accepted.
Issue / escalation object	Trigger conditions, escalation path, remediation owner, SLA, renewal or closure outcome	Moves accepted risk into formal monitoring, escalation, and closure.

Website positioning

Use this artifact as a downloadable governance-layer risk acceptance template and as a future guided skill inside Lapemo. The website version should explain that accepted risk is not a passive exception. It is an operating-model decision that must be owned, evidenced, time-bound, monitored, and connected to AI-safe governance.

Reusable knowledge object structure

Layer	Reusable asset
Human guide	Plain-English explanation, risk categories, register template, workflow, and scale-specific versions.
Downloadable template	DOCX and PDF for governance workshops, risk reviews, executive briefings, and website downloads.
Machine-readable schema	JSON object with required fields, scoring logic, prompts, validation rules, and Lapemo mappings.
Guided skill	Future Lapemo workflow that asks risk acceptance questions, scores readiness, validates evidence, checks AI boundaries, and recommends closure or escalation.