DOCUMENT 1 OF 2 · 10 PAGES

AI Investment Risk Register Template

Enterprise-grade risk register for AI initiatives — taxonomy, identification methodology, scoring, mitigation patterns, review cadence and named owner sign-off.

Audience · Chief Risk Officer, programme owner, model risk, sponsor, audit
Companion documents · Liability Map & Scenario Catalogue
Use as · Audit-grade artefact requiring named human sign-off
Version · v1.0 · 2026-05

§ 1

AI risk taxonomy

Use this taxonomy to classify every risk in the register. The aim is consistent treatment across initiatives and clean aggregation at portfolio level.

Category	Examples
Performance / accuracy	Below-target metric; subgroup regression; concept drift; cold start
Fairness / equity	Disparate impact; disparate treatment; representational harm
Privacy	Unlawful basis; consent failure; PII leakage; cross-border transfer
Security	Endpoint compromise; supply-chain; model inversion; prompt injection
Reliability / availability	Latency; outage; degraded mode without fallback
Safety	Patient harm; environmental release; financial loss to consumer
Regulatory / legal	EU AI Act non-conformity; SR 11-7 / HIPAA / FDA finding; consumer-protection action
Reputational / trust	Public incident; media; community backlash
Operational	HITL operator overload; override pattern; training failure
Financial	Cost overrun; vendor price shock; benefit shortfall
Vendor / third-party	Model deprecation; sub-processor incident; exit friction
Strategic	Mis-fit to strategy; opportunity cost; competitive displacement

§ 2

Risk identification methodology

- **Top-down** · Sponsor + risk identify category-level risks for the initiative.
- **Bottom-up** · Engineering / data / ops walk through the architecture and surface concrete failure modes.
- **Pre-mortem** · "Imagine we are 18 months in and this has gone badly. What did we miss?" Capture top answers.
- **Adversarial** · Red-team / model-attack walk-through (especially for GenAI / agentic systems).
- **External** · Lessons from peer incidents, vendor advisories, regulator findings.

Frequency

- Initial pass at G0; refined at G1; signed at G2; reviewed annually at G3.
- Re-run on any material change (architecture, vendor, scope, regulation).

§ 3

Scoring & tiering

Dimension	1	2	3	4	5
Likelihood	Remote	Unlikely	Possible	Likely	Almost certain
Impact · financial	< $0.1M	$0.1–1M	$1–10M	$10–100M	> $100M
Impact · regulatory	Internal note	Letter	Finding	MRA / Consent	Enforcement
Impact · safety	Inconvenience	Minor harm	Material harm	Serious harm	Catastrophic
Impact · reputational	Internal	Customer-segment	Local press	National press	Global / political

Score = max(impact dimensions) × likelihood. Tier: ≥ 15 = Critical; 9–14 = High; 4–8 = Medium; ≤ 3 = Low.

Heat-map literacy. A heat-map is a communication tool, not a decision tool. Critical & High items must have a named owner, mitigation and target residual.

§ 4

Register template (with examples)

ID	Category	Risk	L	I	Score	Owner	Mitigation	Residual	Review
R-001	Privacy	Consent withdrawal not propagating to feature store	3	4	12	DPO	24 h SLA + CI test + DSR queue audit	Med	Quarterly
R-002	Performance	Offline / online feature skew	4	3	12	ML lead	Feature store with parity tests	Low	Quarterly
R-003	Fairness	Subgroup performance gap on protected class	3	4	12	Fair-lending	AIR monitor + LDA search	Med	Quarterly
R-004	Vendor	Foundation-model deprecation by vendor	3	3	9	Procurement	Multi-vendor optionality; exit plan	Low	Annual
R-005	Security	Prompt-injection on agent tool calls	3	4	12	SecOps	Tool permissioning; output filter; sandbox	Med	Quarterly
R-006	Reg / EU AI Act	Mis-classified as limited-risk; later upgraded to high-risk	2	4	8	AI governance	EU AI Act questionnaire on every release	Low	On change
R-007	Operational	HITL operator throughput overwhelmed at peak	3	3	9	Operations	Surge capacity + queue prioritisation + fallback	Low	Quarterly
R-008	Financial	Vendor / API spend exceeds envelope	3	3	9	Finance	Cap; tiered fallback; monthly review	Low	Monthly
R-009	Reputational	Material adverse incident	2	4	8	Comms + risk	Incident runbook + pre-drafted	Low	Annual

ID	Category	Risk	L	I	Score	Owner	Mitigation	Residual	Review
		becomes public					holding statement		
R-010	Strategic	Initiative becomes obsolete due to vendor / OSS model	2	3	6	Programme	Quarterly horizon scan; portability via canonical schema	Low	Annual

§ 5

Mitigation patterns (library)

Pattern	What it does	When to apply
Permanent holdout	Reconciles claimed lift to truth	Always
Canary + auto-rollback	Limits blast radius of new versions	Every production AI
Kill-switch	Disables AI surface with one flag flip	Operationally critical
Human-in-the-loop	Final action under human authority	Customer / clinical / safety / capital decisions
Diversity / margin floor	Caps over-personalisation impact	Recommender / pricing
Fairness guardrail	Pause on protected-class disparity	Any model used in regulated decisions
Cap & tiered fallback	Bound vendor / cost exposure	Foundation-model dependent
Independent validation	Effective challenge against 1st line	SR 11-7 / clinical / EU AI Act high-risk
Annual re-DPIA / re-validation	Catches latent regulatory drift	Always
Multi-vendor optionality	Reduces lock-in & deprecation risk	High-spend FM / API

§ 6

Review cadence

Forum	Cadence	Scope
Programme stand-up	Weekly	New & Critical risks; open mitigations
Risk & control board	Monthly	Material risks across portfolio; aggregated heat-map
Sponsor review	Quarterly	Critical & High; residual trend
Annual review	Annual	Full re-pass; tied to G3 review
Audit committee briefing	Annual or on event	Heat-map + top 5; companion briefing pack
Board	Annual or on event	Portfolio-level AI risk posture

§ 7

Aggregation & portfolio view

Initiative-level registers roll up to a portfolio-level view, useful for board and audit committees.

Initiative	Critical	High	Medium	Top open risk	Owner
Retail personalisation	0	2	5	R-002 feature skew	ML lead
Credit scoring v3	0	1	4	R-003 fairness	Fair-lending
Triage copilot	0	1	3	Clinical safety	CMO
Predictive maintenance	0	0	2	Sensor drift	Reliability

Portfolio-level numbers are recomputed monthly. A new Critical or sustained High at portfolio level triggers a CRO briefing.

§ 8

Connection to other registers

The AI register is not a separate animal. It rolls into enterprise risk and connects to several sibling registers.

- Operational risk register (people, process, systems, external)
- Information-security risk register
- Privacy risk register (records of processing)
- Vendor / third-party risk register
- Compliance risk register (regulatory)
- Strategic / business-line risk register

Where the same risk lives in multiple registers, the AI register stays the source of truth for the AI-specific framing; cross-references are maintained.

§ 9

Sample remediation log

Risk #	Action	Owner	Due	Status	Verified
R-001	Add consent-withdrawal CI test	Eng			
R-003	Implement LDA search and document	Fair-lending			
R-005	Tool-call allow-list and sandbox	SecOps			

§ 10

Sign-off & attestation

Sign-off block

Role	Name	Date	Signature
Chief Risk Officer (or delegate)			
Programme sponsor			
Programme owner			
Model risk lead			
Privacy / DPO			
Security			

The register is signed at each gate (companion gating workbook) and at the annual review. Critical risks in residual state require sponsor + CRO co-sign with documented acceptance.