

SIMULATED CASE STUDY

AI Risk Assessment: Policy and Procedure Document Analyser

AI governance

Project	AI governance and use-case risk assessment (Project 4)
Company	Northstar Financial Technologies (fictional)
Status	Simulated; no certification or conformance claimed
As of	2026-09-30

This document is part of a portfolio exercise. Northstar Financial Technologies, its people, systems, records and findings are invented. No real audit was performed, no assurance is given and nothing here is a certification or compliance claim.

1. Use case

Use case	AI-02 Policy and Procedure Document Analyser
Purpose	Summarise internal policy and procedure documents and answer questions about them for employees
Business owner	Chief Financial Officer
Users	Employees in Risk, Compliance, Finance and Legal (about 70)
Data	Internal policies, procedures, contracts and board papers that may include confidential and personal information
Affected parties	Employees; indirectly customers and counterparties named in documents
Provider	Same third-party LLM provider, document-upload tier (vendor V-03, Project 3 residual risk High)
Proposed oversight	Users are expected to verify summaries against source documents (no technical control proposed)
Overall rating: inherent / as proposed / after gates	High / High / Medium
Decision	Not approved: must not proceed

2. How ratings work

Each of eleven areas is rated Low, Medium or High three times: inherent (no controls), as proposed in the intake, and after the gates are met. The overall rating is High if three or more areas are High, Medium if any area is High or three or more are Medium, otherwise Low. The scale is ordinal judgement; the rationale column is the substance.

3. Assessment by area

Area	Inherent	As proposed	After gates	Rationale	Controls required
Intended purpose and affected people	Medium	Medium	Medium	Employees in Risk, Compliance, Finance and Legal rely on summaries; customers and counterparties named in documents are indirectly affected.	Defined permitted document types; summaries marked as AI-generated.
Business criticality and fallback	Low	Low	Low	A convenience tool; staff can read the source documents.	None beyond fallback to reading the document.
Data categories, privacy and confidentiality	High	High	Medium	Policies, contracts and board papers can contain confidential and personal information and legally privileged content. The document-upload tier's retention and training terms are unknown.	Classification gate: only Internal and permitted Confidential documents; exclude board papers, privileged and Restricted documents; contract confirms no training use and defined retention for document upload; privacy impact assessment.
Output reliability and hallucination	High	High	Medium	Summaries can omit or distort a clause that matters for a compliance decision.	Every answer cites the source passage; users must open the cited passage for any decision; periodic accuracy test on a policy set.
Human oversight	High	High	Medium	The intake relies on users to verify summaries 'with no technical control proposed', which is not a control.	Technical traceability (source passage shown), required acknowledgement before copying output into a record, training.
Security: prompt injection and data leakage	High	High	Medium	Uploaded documents can contain hidden instructions (indirect prompt injection), and without per-user access control one user could see another's documents.	Per-user document isolation; users can upload only what they may access; scanning of documents; no tool actions for the model; injection tests.
Supplier and subprocessor dependency	High	High	Medium	The same vendor, but the document-upload tier is outside the evidence Northstar holds. Project 3 restricts V-03 from this use.	Close all Project 3 conditions and obtain written terms for the upload tier, or choose a different provider or a self-hosted option.
Logging and traceability	Medium	High	Low	No record would show who uploaded which document or who saw which output.	Log uploader, document id and classification, query and answer; review access monthly.
Change management (model and vendor changes)	Medium	Medium	Low	Vendor changes to upload processing could change retention without notice.	Notice-of-change clause; re-assess on any change.
Accessibility and fairness	Low	Low	Low	Few affected individuals; no decisions about people are made.	Prohibit use for employee or customer decisions.
Incident response	High	High	Medium	Exposure of uploaded confidential documents to a third party is hard to reverse and to scope.	Adopt the AI incident playbook; contract right to deletion confirmation; evidence of what was uploaded from logs.

4. Decision and gates

Not approved: must not proceed. No use of the document-upload tier or any external AI service with Northstar internal documents until gates AI-02-G01 to G08 are evidenced and the use case is resubmitted. Interim: employees must not upload Northstar documents to Lumenvale.

The documents include confidential, personal and privileged content, the vendor's upload-tier terms are unknown, and the intake relies on users to catch errors with no technical control. The residual risk is High on privacy, oversight, security and supplier dependency together, and the business value does not justify the exposure while those are

open.

Gate	Requirement	Owner	Target	Status
AI-02-G01	Written terms for the document-upload tier: retention, no model training, subprocessors, incident notification; or an alternative provider	Chief Financial Officer	2027-01-31	Blocked by 3 open vendor issue(s) in Project 3
AI-02-G02	Classification gate that blocks upload of board papers, privileged and Restricted documents	Head of Platform Engineering	2027-01-31	Not started
AI-02-G03	Per-user document isolation and upload limited to documents the user may access	Head of Platform Engineering	2027-01-31	Not started
AI-02-G04	Technical traceability: every answer shows the source passage and requires acknowledgement before reuse in a record	Chief Financial Officer	2027-01-31	Not started
AI-02-G05	Privacy impact assessment and legal review of privilege and confidentiality risks	Privacy Officer	2027-01-31	Not started
AI-02-G06	Accuracy evaluation on a representative set of policy documents with documented thresholds	Chief Financial Officer	2027-02-28	Not started
AI-02-G07	Upload and access logging with monthly access review	IT Security Manager	2027-02-28	Not started
AI-02-G08	Monitoring plan and AI incident playbook adopted and rehearsed	AI Governance Committee	2027-02-28	Not started
The residual rating after gates assumes every gate is closed with evidence. Until then the as-proposed rating applies. Acceptance authority: Chief Risk Officer (executive acceptance) for the as-proposed level; AI Governance Committee after gates.				