

SIMULATED CASE STUDY

# AI Governance Policy

Lightweight policy for Northstar

|         |                                                        |
|---------|--------------------------------------------------------|
| 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. Purpose and scope

This lightweight policy sets how Northstar decides whether an AI system may be used, under what conditions, and how it is watched afterwards. It covers any AI service or model used for Northstar work, whether bought, embedded in another product or built in-house. It is a simulated policy for a fictional company.

# 2. Principles

- AI output is advisory. A named person is accountable for every decision or message that relies on it.
- No use without intake, risk assessment and vendor due diligence.
- Use the least data that does the job; Restricted data never goes to an external AI service.
- Approval is conditional on evidence, not on promises. Gates are closed only with evidence.
- Anything that cannot be switched off quickly is not approved.

# 3. Roles and responsibilities

| Role                                         | Responsibility                                                                                             |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------|
| Business owner                               | Accountable for the use case, its benefits and its residual risk within appetite; proposes human oversight |
| AI Governance Committee                      | Approves use cases and gates; accepts Medium residual risk; reviews monitoring                             |
| Chief Risk Officer                           | Accepts High residual risk; owns this policy                                                               |
| Privacy Officer                              | Privacy impact assessment; advises on notification decisions                                               |
| IT Security Manager                          | Security testing, logging review, incident lead                                                            |
| Legal                                        | Contract terms, privilege and regulatory advice                                                            |
| Head of Procurement                          | Vendor due diligence and contract conditions (Project 3)                                                   |
| Model or system owner (Platform Engineering) | Configuration, change control, logs, switch-off                                                            |

# 4. Lifecycle and approval gates

| Gate                         | What must be true                                                                                                                                               |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| G0 Intake                    | Business owner submits purpose, users, data, vendor, proposed oversight                                                                                         |
| G1 Risk assessment           | Assessed across eleven risk areas with inherent and residual ratings and rationale                                                                              |
| G2 Specialist review         | Privacy, security, legal and procurement reviews complete; vendor due diligence (Project 3 method) done                                                         |
| G3 Pre-deployment validation | Evaluation on a representative set with thresholds; security tests; oversight enforced technically                                                              |
| G4 Decision                  | Approve, conditionally approve (with gates), or do not proceed. High residual risk is accepted by the Chief Risk Officer; Medium by the AI Governance Committee |
| G5 Operation and monitoring  | Monitoring measures live; periodic review; suspension triggers enforced                                                                                         |
| G6 Retirement                | Data deleted, access removed, vendor exit confirmed, inventory updated                                                                                          |

# 5. Prohibited and restricted use

## Prohibited

- Automated decisions with legal or similarly significant effect on a person without a human decision

- Sending AI-generated output to customers without human review
- Entering Restricted data (card data, credentials, cryptographic keys) into any external AI service
- Employee surveillance or profiling using AI
- Use of an AI service that has not completed the intake and vendor due diligence
- Using AI output as the sole basis of a regulatory filing, audit conclusion or legal position

### Restricted (allowed only on the conditions stated)

- Confidential data may be used only with an approved vendor tier whose retention and training terms are signed
- Customer-facing drafting is allowed only with technically enforced human approval
- Document analysis is limited to Internal and permitted Confidential documents; board papers and privileged documents are excluded
- Audit or compliance analysis by AI is advisory only and a person decides (see Project 5)

## 6. Approved data classes

| Class        | Rule                                                                                        |
|--------------|---------------------------------------------------------------------------------------------|
| Public       | Allowed with an approved tool                                                               |
| Internal     | Allowed with an approved tool                                                               |
| Confidential | Allowed only for an approved use case, approved vendor tier and minimisation where possible |
| Restricted   | Not allowed in external AI services                                                         |

## 7. Human review

Customer-facing output requires a person to approve each item and the tool must prevent sending without approval. Internal output that informs a decision must show its sources and the user must check them. Approval counts as review only if reviewers are trained and their edit and override rates are monitored, because routine approval is a known weakness (automation bias).

## 8. Escalation and exceptions

A use outside this policy needs a written exception from the Chief Risk Officer, with a reason, compensating controls, an end date of no more than 90 days and a record in the inventory. Triggers for escalation: a privacy or security incident, a High residual risk without acceptance, a vendor term change, or use beyond the approved scope.

## 9. Monitoring and review

| When                 | Review                                                                                                  |
|----------------------|---------------------------------------------------------------------------------------------------------|
| Pilot period (AI-01) | Monthly review by the AI Governance Committee for six months, then quarterly                            |
| Steady state         | Quarterly review of every approved use case; annual full reassessment                                   |
| Event-driven         | Within 10 business days of a suspension, incident, vendor change, model change or material scope change |
| Policy               | Reviewed annually by the Chief Risk Officer                                                             |

Suspend a use case on: A confirmed privacy or security incident; Material error rate above the red threshold; A vendor change not re-tested; Use outside the approved scope or data classes; An open High vendor issue overdue

more than 30 days.

## 10. Relationship to frameworks

The risk structure follows the four functions of the NIST AI Risk Management Framework 1.0: Govern (this policy, roles, inventory, third-party policies), Map (intake, purpose and impact), Measure (evaluation and monitoring measures) and Manage (treatment, gates, incident response and retirement). The mapping uses subcategories verified in the NIST document and is a judgement about support, not conformance. ISO/IEC 42001:2023 is a management system standard for AI; this programme has the shape of a lightweight management system, but no clause-level mapping was made, and Northstar makes no claim of alignment, readiness or certification.

## 11. Limitations

Simulated policy for a fictional company, not legal advice. Thresholds in the monitoring plan are proposals to be calibrated. NIST has said the AI RMF is under revision, so references should be rechecked.