SENTRAI | GOVERNANCE

Northwind Finance (sample)
ORGANISATION

# AI MODEL GOVERNANCE SUMMARY

IT HELPDESK ANSWER ASSISTANT | NORTHWIND FINANCE (SAMPLE) · LENDING AND OPERATIONS AI

DRAFT (LIVE)

## EXECUTIVE GOVERNANCE SNAPSHOT

This page summarises the card and highlights where evidence or validation is still needed. The four scores are separate measures and are never combined into one number.

| 67 / 88 | 0 / 14 | 0 / 16 | 1 / 2 |
|---|---|---|---|
| DOCUMENTATION | EVIDENCE | CONTROL EFFECTIVENESS | REGULATORY ASSESSMENT |
| 76% complete | 0% complete | 0% complete | 50% complete |

## RECORD AT A GLANCE

| Attribute | Recorded value | Attribute | Recorded value |
|---|---|---|---|
| **Model / record name** | IT Helpdesk Answer Assistant | **Record type** | Ai system |
| **Component lifecycle** | In test | **Model lifecycle** | Validation |
| **Deployment status** | Deployed non production | **Operational status** | Active |
| **Inherent risk tier** | Low | **Residual risk tier** | Low (Assessed) |
| **EU AI Act classification** | — | **Korea high-impact** | — |
| **Workflow state** | Evidence collection | **Next review** | — (Not scheduled) |

## FRAMEWORK STATUS

| Framework (enabled for the organisation) | Automated clause check |
|---|---|
| NIST AI RMF 1.0 | 4 open gaps |
| ISO/IEC 42001:2023 | No open gaps (6 met) |

*Automated checks and registry mappings; not an assessment, audit or certification.*

## RISK POSTURE

Inherent Low, residual Low (Assessed). Risk engine: score 4 of 23, score tier Low, Low with floors. A person sets the tiers; the engine explains and checks them.

## KEY GOVERNANCE ATTENTION POINTS

| Priority | Finding | Why it matters |
|---|---|---|
| **High** | 23 items block a final card | Provider may train on inputs is UNKNOWN; find out and record it (required because R_THIRD_PARTY fired). Sub-processors is UNKNOWN; find out and record it (required because R_THIRD_PARTY fired). Data processing agreements is required (required because R_THIRD_PARTY fired). See section 1. |
| **High** | Risk tier needs reassessment | LOW is declared without evaluation evidence. Human oversight or limited autonomy alone does not establish low risk. |
| **High** | Evaluation evidence absent | No evaluation evidence, red-team date or independent validation is recorded. |
| **Medium** | Data and dependencies incomplete | Datasets and dependencies are blank, so data flows and personal-data processing cannot be confirmed. |
| **Medium** | NIST AI RMF 1.0: 4 open gaps | Mandatory clauses not yet met: GOVERN 1.5, MEASURE 2.7, MEASURE 2.4, MEASURE 2.10. |
| **Low** | No audit scheduled | Next audit date is not set. |

DRAFT

# 1. Finalisation gate

| # | Blocking a final card |
|---|---|
| 1 | Provider may train on inputs is UNKNOWN; find out and record it (required because R_THIRD_PARTY fired). |
| 2 | Sub-processors is UNKNOWN; find out and record it (required because R_THIRD_PARTY fired). |
| 3 | Data processing agreements is required (required because R_THIRD_PARTY fired). |
| 4 | Provider certifications is required (required because R_THIRD_PARTY fired). |
| 5 | Third-party risk assessment is required (required because R_THIRD_PARTY fired). |
| 6 | Model update policy is required (required because R_THIRD_PARTY fired). |
| 7 | User disclosure is required (required because R_GENERATIVE fired). |
| 8 | AI-content labelling is required (required because R_GENERATIVE fired). |
| 9 | Harmful-output evaluation is required (required because R_GENERATIVE fired). |
| 10 | Prompt-injection testing is required (required because R_GENERATIVE fired). |
| 11 | Next revalidation date is UNKNOWN; find out and record it. |
| 12 | Incident response is required. |
| 13 | Moderation is required (required because R_GENERATIVE fired). |
| 14 | Output filtering is required (required because R_GENERATIVE fired). |
| 15 | Guardrails is required (required because R_GENERATIVE fired). |
| 16 | Safety testing is required (required because R_GENERATIVE fired). |
| 17 | Legal review date is required. |
| 18 | Exit strategy is required. |
| 19 | Responsibility split is required. |
| 20 | Regulatory notification is required (required because R_HAS_JURISDICTION fired). |
| 21 | Suspension criteria is required (required because R_DEPLOYED fired). |
| 22 | Model lifecycle is VALIDATION but operational status is ACTIVE. A record that has not been deployed cannot be operating. |
| 23 | The component is IN_TEST but operational status is ACTIVE. An unreleased component cannot be operating. |

## Warnings

| # | Warning |
|---|---|
| 1 | Developer is recommended. |
| 2 | Release date is recommended. |
| 3 | Licence is recommended. |
| 4 | Deployment identifier is recommended. |
| 5 | Business owner is recommended. |
| 6 | Registration date is recommended. |
| 7 | Last updated is recommended. |
| 8 | Observed use is recommended. |
| 9 | Target population is recommended. |
| 10 | Tasks is recommended. |
| 11 | Languages is recommended. |
| 12 | Operating environment is recommended. |
| 13 | Expected outputs is recommended. |
| 14 | Foreseeable misuse is recommended. |

DRAFT

| # | Warning |
|---|---|
| 15 | Expected performance is recommended. |
| 16 | Usage restrictions is recommended. |
| 17 | Architecture description is recommended. |
| 18 | Input modalities is recommended. |
| 19 | Output modalities is recommended. |
| 20 | Formats is recommended. |
| 21 | Fine-tuning is recommended. |
| 22 | Runtime is recommended. |
| 23 | API integration is recommended. |
| 24 | Hosting and deployment architecture is recommended. |
| 25 | Serving technology is recommended. |
| 26 | Training data sources is recommended. |
| 27 | Training data period is recommended. |
| 28 | Data ownership is recommended. |
| 29 | Data lineage is recommended. |
| 30 | Data quality is recommended. |
| 31 | Representativeness is recommended. |
| 32 | Copyright and IP is recommended. |
| 33 | Data-poisoning assessment is recommended. |
| 34 | Synthetic data is recommended. |
| 35 | Foundation, parent or fine-tuned-from model is recommended. |
| 36 | Embedding dependencies is recommended. |
| 37 | External APIs is recommended. |
| 38 | Third-party services is recommended. |
| 39 | Open-source libraries and versions is recommended. |
| 40 | SBOM or AI-BOM reference is recommended. |
| 41 | End-of-life and deprecation policy is recommended. |
| 42 | Organisations reached is recommended. |
| 43 | Tenants reached is recommended. |
| 44 | Existing controls is recommended. |
| 45 | Risk acceptance is recommended. |
| 46 | Societal impact is recommended. |
| 47 | Financial impact is recommended. |
| 48 | Employment impact is recommended. |
| 49 | Health and safety impact is recommended. |
| 50 | Environmental impact is recommended. |
| 51 | Stakeholder consultation is recommended. |
| 52 | Representativeness is recommended. |
| 53 | Mitigation is recommended. |
| 54 | Independent evidence is recommended. |
| 55 | Interpretability is recommended. |

DRAFT

| # | Warning |
|---|---|
| 56 | Explainability limitations is recommended. |
| 57 | Traceability is recommended. |
| 58 | Notification mechanisms is recommended. |
| 59 | Review frequency is recommended. |
| 60 | Intervention and override is recommended. |
| 61 | Escalation is recommended. |
| 62 | Reviewer competency is recommended. |
| 63 | AI literacy training is recommended. |
| 64 | Automation boundaries is recommended. |
| 65 | Tenant notification is recommended. |
| 66 | Cross-tenant leakage testing is recommended. |
| 67 | Access control is recommended. |
| 68 | Encryption at rest is recommended. |
| 69 | Encryption in transit is recommended. |
| 70 | Model-extraction protection is recommended. |
| 71 | Adversarial resistance is recommended. |
| 72 | Poisoning protection is recommended. |
| 73 | Supply-chain security is recommended. |
| 74 | Vulnerability management is recommended. |
| 75 | Rate limiting is recommended. |
| 76 | Abuse prevention is recommended. |
| 77 | Disaster recovery is recommended. |
| 78 | Rollback is recommended. |
| 79 | Business continuity is recommended. |
| 80 | Security assessment evidence is recommended. |
| 81 | Safety policies is recommended. |
| 82 | Harm-prevention controls is recommended. |
| 83 | Misuse prevention is recommended. |
| 84 | Dangerous-capability assessment is recommended. |
| 85 | Robustness safeguards is recommended. |
| 86 | Safety incidents is recommended. |
| 87 | Residual safety risk is recommended. |
| 88 | Other regulatory classifications is recommended. |
| 89 | Exemptions is recommended. |
| 90 | Next regulatory review is recommended. |
| 91 | Audit evidence is recommended. |
| 92 | Certifications is recommended. |
| 93 | Sub-processors is recommended. |
| 94 | Incident history is recommended. |
| 95 | Concentration risk is recommended. |
| 96 | Environment is recommended. |

DRAFT

| # | Warning |
|---|---|
| 97 | Deployment date is recommended. |
| 98 | Deployment approval is recommended. |
| 99 | Change history is recommended. |
| 100 | Change-management process is recommended. |
| 101 | Decommissioning date is recommended. |
| 102 | Open incidents is recommended. |
| 103 | Incident severity and classification is recommended. |
| 104 | Affected users is recommended. |
| 105 | Root cause is recommended. |
| 106 | Corrective action is recommended. |
| 107 | Remediation status is recommended. |
| 108 | Inference energy is recommended. |
| 109 | Carbon emissions is recommended. |
| 110 | Hosting region is recommended. |
| 111 | Efficiency measures is recommended. |
| 112 | Assessment date is recommended. |

# 2. Card sections

## 1. Identity and metadata (15/15 required complete)

| Field | Value | Status |
|---|---|---|
| **Record type *** | AI_SYSTEM | Populated · required · entered by a person |
| **Model or record identifier *** | NW-AI-0012 | Populated · required · entered by a person |
| **Name *** | IT Helpdesk Answer Assistant | Populated · required · entered by a person |
| **Provider *** | Asterion AI (hosted language model API) | Populated · required · entered by a person |
| **Developer** | unknown | Populated · recommended · imported from v1 · awaiting acceptance |
| **Model category *** | LLM | Populated · required · entered by a person |
| **Version *** | 0.9.0 | Populated · required · entered by a person |
| **Release date** | — | Not assessed · recommended |
| **Licence** | — | Not assessed · recommended |
| **Access type** | API | Populated · recommended · entered by a person |
| **Hosting type *** | PROVIDER_HOSTED | Populated · required · entered by a person |
| **Hosting region *** | India (Mumbai) | Populated · required · entered by a person |
| **Deployment identifier** | — | Not assessed · recommended |
| **Component lifecycle *** | IN_TEST | Populated · required · entered by a person |
| **Model lifecycle *** | VALIDATION | Populated · required · entered by a person |
| **Deployment status *** | DEPLOYED_NON_PRODUCTION | Populated · required · entered by a person |
| **Operational status *** | ACTIVE | Populated · required · entered by a person |
| **Accountable owner *** | 0e4c3550-0634-4acb-8262-84b52454853f<br>Refs: nist-ai-rmf-1.0:GOVERN-2, iso-42001-2023:A.3 | Populated · required · entered by a person |
| **Technical owner *** | 1a1a23d1-7cee-4fa4-877b-639306d9c76b | Populated · required · entered by a person |
| **Business owner** | — | Not assessed · recommended |

DRAFT

| Field | Value | Status |
|---|---|---|
| Registration date | 2026-11-27 | Populated · recommended · imported from v1 · awaiting acceptance |
| Last updated | — | Not assessed · recommended |
| Next review date * | 2027-02-28 | Populated · required · entered by a person |

## 2. Intended use and scope (12/12 required complete)

| Field | Value | Status |
|---|---|---|
| Business justification * | Staff raise about 2,000 routine IT questions a month. Answers drawn from the IT knowledge base could resolve half of them without a ticket. | Populated · required · entered by a person |
| Purpose * | Answers employee IT questions (passwords, VPN, printers, software requests) from the internal IT knowledge base.<br>Refs: nist-ai-rmf-1.0:MAP-1, iso-42001-2023:A.6 | Populated · required · entered by a person |
| Intended use * | Employees ask a question in the intranet chat; the assistant answers with a link to the knowledge-base article or offers to raise a ticket. | Populated · required · entered by a person |
| Permitted use * | Internal staff, IT topics only. | Populated · required · entered by a person |
| Observed use | — | Not assessed · recommended |
| Intended users * | Employees | Populated · required · entered by a person |
| Affected segments * | Employees<br>Refs: nist-ai-rmf-1.0:MAP-5 | Populated · required · entered by a person |
| Target population | — | Not assessed · recommended |
| Tasks | — | Not assessed · recommended |
| Capabilities * | Answers IT questions with a source link; Offers to raise a ticket | Populated · required · entered by a person |
| Languages | — | Not assessed · recommended |
| Geographic scope * | India | Populated · required · entered by a person |
| Operating environment | — | Not assessed · recommended |
| Expected outputs | — | Not assessed · recommended |
| Decision impact * | INFORMATIONAL | Populated · required · entered by a person |
| Out-of-scope uses * | HR, payroll or customer questions | Populated · required · entered by a person |
| Prohibited uses * | Resetting credentials or changing access | Populated · required · entered by a person |
| Foreseeable misuse | —<br>Refs: eu-ai-act-2024:ART-9 | Not assessed · recommended |
| Known limitations * | Only knows what is in the knowledge base; English only | Populated · required · entered by a person |
| Expected performance | — | Not assessed · recommended |
| Usage restrictions | — | Not assessed · recommended |

## 3. Technical specifications (2/2 required complete)

| Field | Value | Status |
|---|---|---|
| Architecture description | — | Not assessed · recommended |
| Input modalities | — | Not assessed · recommended |
| Output modalities | — | Not assessed · recommended |
| Formats | — | Not assessed · recommended |
| Inference configuration * | Temperature 0.2, maximum 600 output tokens, no tools. | Populated · required · entered by a person |
| Prompts and system prompt * | (sensitive: not printed) | Populated · required · entered by a person |

DRAFT

| Field | Value | Status |
|---|---|---|
| Fine-tuning | — | Not assessed · recommended |
| Runtime | — | Not assessed · recommended |
| API integration | — | Not assessed · recommended |
| Hosting and deployment architecture | — | Not assessed · recommended |
| Serving technology | — | Not assessed · recommended |

## 4. Data governance and provenance (2/3 required complete)

| Field | Value | Status |
|---|---|---|
| Training data sources | —<br>Refs: eu-ai-act-2024:ART-10, iso-42001-2023:A.7 | Not assessed · recommended |
| Training data period | — | Not assessed · recommended |
| Data ownership | — | Not assessed · recommended |
| Data lineage | — | Not assessed · recommended |
| Data classification * | INTERNAL | Populated · required · entered by a person |
| Data quality | — | Not assessed · recommended |
| Representativeness | — | Not assessed · recommended |
| Personal data processed * | NO<br>Refs: gdpr:ART-6, in-dpdp-2023:NOTICE | Populated · required · entered by a person |
| Copyright and IP | —<br>Refs: nist-ai-600-1:INTELLECTUAL_PROPERTY | Not assessed · recommended |
| Data-poisoning assessment | — | Not assessed · recommended |
| Synthetic data | — | Not assessed · recommended |
| Provider may train on inputs * | UNKNOWN<br>Refs: nist-ai-600-1:DATA_PRIVACY | Unknown · required · entered by a person |

## 5. Dependencies and supply chain (3/8 required complete)

| Field | Value | Status |
|---|---|---|
| Third-party provider used * | YES | Populated · required · entered by a person |
| Foundation, parent or fine-tuned-from model | — | Not assessed · recommended |
| Embedding dependencies | — | Not assessed · recommended |
| External APIs | — | Not assessed · recommended |
| Third-party services | — | Not assessed · recommended |
| Sub-processors * | UNKNOWN | Unknown · required · entered by a person |
| Open-source libraries and versions | — | Not assessed · recommended |
| SBOM or AI-BOM reference | — | Not assessed · recommended |
| Provider documentation * | https://docs.northwind.example/ai/it-helpdesk-assistant/provider-docs<br>Evidence: provider_documentation https://docs.northwind.example/ai/it-helpdesk-assistant/provider-docs | Populated · required · entered by a person |
| SLAs * | 99.9% monthly availability | Populated · required · entered by a person |
| Data processing agreements * | — | Not assessed · required |
| Provider certifications * | — | Not assessed · required |

DRAFT

| Field | Value | Status |
|---|---|---|
| Third-party risk assessment * | — Refs: iso-42001-2023:A.10, nist-ai-600-1:VALUE_CHAIN | Not assessed · required |
| Model update policy * | — | Not assessed · required |
| End-of-life and deprecation policy | — | Not assessed · recommended |

## 6. Risk classification (10/10 required complete)

| Field | Value | Status |
|---|---|---|
| Impact severity * | MINOR | Populated · required · entered by a person |
| Likelihood * | POSSIBLE | Populated · required · entered by a person |
| Users reached * | 4200 | Populated · required · entered by a person |
| Organisations reached | — | Not assessed · recommended |
| Tenant reach * | ONE_TENANT | Populated · required · entered by a person |
| Tenants reached | — | Not assessed · recommended |
| Autonomy level * | NONE | Populated · required · entered by a person |
| Data sensitivity * | LOW | Populated · required · entered by a person |
| Reversibility * | EASILY_REVERSED | Populated · required · entered by a person |
| Existing controls | — | Not assessed · recommended |
| Inherent risk tier * | LOW Refs: nist-ai-rmf-1.0:MAP-4, iso-23894-2023:RISK_PROCESS | Populated · required · entered by a person |
| Residual risk tier * | LOW Refs: nist-ai-rmf-1.0:MANAGE-1 | Populated · required · entered by a person |
| Risk rationale * | Wrong answers are easy to spot and cost a little time; no personal data and no actions. Residual risk equals inherent risk until the evaluation is complete. | Populated · required · entered by a person |
| Risk acceptance | — | Not assessed · recommended |

## 7. Impact assessment (0/0 required complete)

| Field | Value | Status |
|---|---|---|
| Societal impact | — | Not assessed · recommended |
| Financial impact | — | Not assessed · recommended |
| Employment impact | — | Not assessed · recommended |
| Health and safety impact | — | Not assessed · recommended |
| Environmental impact | — | Not assessed · recommended |
| Stakeholder consultation | — | Not assessed · recommended |

## 8. Fairness and bias (0/0 required complete)

| Field | Value | Status |
|---|---|---|
| Representativeness | — | Not assessed · recommended |
| Mitigation | — | Not assessed · recommended |
| Independent evidence | — | Not assessed · recommended |

## 9. Explainability and transparency (0/2 required complete)

| Field | Value | Status |
|---|---|---|
| Interpretability | — | Not assessed · recommended |
| Explainability limitations | — | Not assessed · recommended |

DRAFT

| Field | Value | Status |
|---|---|---|
| **Traceability** | —<br>Refs: eu-ai-act-2024:ART-12 | Not assessed · recommended |
| **User disclosure *** | —<br>Refs: eu-ai-act-2024:ART-50, kr-ai-basic-act:USER_NOTICE | Not assessed · required |
| **AI-content labelling *** | —<br>Refs: eu-ai-act-2024:ART-50, kr-ai-basic-act:OUTPUT_LABEL, in-it-rules-2026:SYNTHETIC_LABEL, cn-ai-labelling:LABEL | Not assessed · required |
| **Notification mechanisms** | — | Not assessed · recommended |

## 10. Human oversight and accountability (2/2 required complete)

| Field | Value | Status |
|---|---|---|
| **Oversight model *** | HUMAN_ON_THE_LOOP<br>Refs: eu-ai-act-2024:ART-14, nist-ai-600-1:HUMAN_AI_CONFIGURATION | Populated · required · entered by a person |
| **Responsible role *** | IT service desk lead (weekly review of answers) | Populated · required · entered by a person |
| **Review frequency** | — | Not assessed · recommended |
| **Intervention and override** | — | Not assessed · recommended |
| **Escalation** | — | Not assessed · recommended |
| **Reviewer competency** | — | Not assessed · recommended |
| **AI literacy training** | —<br>Refs: eu-ai-act-2024:ART-4 | Not assessed · recommended |
| **Automation boundaries** | — | Not assessed · recommended |
| **Tenant notification** | — | Not assessed · recommended |

## 11. Evaluation and validation (2/5 required complete)

| Field | Value | Status |
|---|---|---|
| **Evaluation records *** | Answer accuracy on 200 past tickets / ACCURACY / PRE_DEPLOYMENT / 2026-11-10 / IT service desk / 88% correct with source / 90% / FAIL / https://docs.northwind.example/ai/it-helpdesk-assistant/tests/accuracy-2026-11 / Add missing VPN and printer articles, then re-test<br>Refs: nist-ai-rmf-1.0:MEASURE-2, eu-ai-act-2024:ART-15 | Populated · required · entered by a person |
| **Hallucination evaluation *** | https://docs.northwind.example/ai/it-helpdesk-assistant/tests/accuracy-2026-11<br>Evidence: evaluation_report https://docs.northwind.example/ai/it-helpdesk-assistant/tests/accuracy-2026-11<br>Refs: nist-ai-600-1:CONFABULATION | Populated · required · entered by a person |
| **Harmful-output evaluation *** | —<br>Refs: nist-ai-600-1:DANGEROUS_CONTENT | Not assessed · required |
| **Prompt-injection testing *** | —<br>Refs: nist-ai-600-1:INFORMATION_SECURITY | Not assessed · required |
| **Cross-tenant leakage testing** | — | Not assessed · recommended |
| **Next revalidation date *** | UNKNOWN | Unknown · required · entered by a person |

## 12. Security and resilience (4/5 required complete)

| Field | Value | Status |
|---|---|---|
| **Authentication *** | Bank single sign-on | Populated · required · entered by a person |
| **Authorisation *** | All employees | Populated · required · entered by a person |
| **Access control** | — | Not assessed · recommended |
| **API-key and BYOK key management *** | (sensitive: not printed) | Populated · required · entered by a person |

DRAFT

| Field | Value | Status |
|---|---|---|
| Encryption at rest | — | Not assessed · recommended |
| Encryption in transit | — | Not assessed · recommended |
| Prompt-injection protection * | Knowledge-base text is passed as data; injection testing not yet done.<br>Refs: nist-ai-600-1:INFORMATION_SECURITY | Populated · required · entered by a person |
| Model-extraction protection | — | Not assessed · recommended |
| Adversarial resistance | — | Not assessed · recommended |
| Poisoning protection | — | Not assessed · recommended |
| Supply-chain security | — | Not assessed · recommended |
| Vulnerability management | — | Not assessed · recommended |
| Rate limiting | — | Not assessed · recommended |
| Abuse prevention | — | Not assessed · recommended |
| Incident response * | — | Not assessed · required |
| Disaster recovery | — | Not assessed · recommended |
| Rollback | — | Not assessed · recommended |
| Business continuity | — | Not assessed · recommended |
| Security assessment evidence | —<br>Refs: eu-ai-act-2024:ART-15 | Not assessed · recommended |

## 13. Safety and responsible AI (0/4 required complete)

| Field | Value | Status |
|---|---|---|
| Safety policies | AI-POL-2026 | Populated · recommended · imported from v1 · awaiting acceptance |
| Harm-prevention controls | — | Not assessed · recommended |
| Moderation * | — | Not assessed · required |
| Output filtering * | — | Not assessed · required |
| Guardrails * | — | Not assessed · required |
| Misuse prevention | — | Not assessed · recommended |
| Dangerous-capability assessment | —<br>Refs: eu-ai-act-2024:ART-55, nist-ai-600-1:CBRN | Not assessed · recommended |
| Robustness safeguards | — | Not assessed · recommended |
| Safety testing * | — | Not assessed · required |
| Safety incidents | — | Not assessed · recommended |
| Residual safety risk | — | Not assessed · recommended |

## 14. Regulatory and legal (5/6 required complete)

| Field | Value | Status |
|---|---|---|
| Jurisdictions * | IN | Populated · required · entered by a person |
| India: acts as an intermediary * | NO | Populated · required · entered by a person |
| Other regulatory classifications | — | Not assessed · recommended |
| Applicable laws * | Information Technology Act 2000 | Populated · required · entered by a person |
| Applicability rationale * | Internal tool, no personal data, India only. | Populated · required · entered by a person |
| Exemptions | — | Not assessed · recommended |

DRAFT

| Field | Value | Status |
|---|---|---|
| **Compliance owner *** | 7bb09eaa-f07e-4c0b-a26c-bc863b729cea | Populated · required · entered by a person |
| **Legal review date *** | — | Not assessed · required |
| **Next regulatory review** | — | Not assessed · recommended |

## 15. Third-party governance (5/7 required complete)

| Field | Value | Status |
|---|---|---|
| **Provider *** | Asterion AI<br>Refs: iso-42001-2023:A.10 | Populated · required · entered by a person |
| **Provider location *** | India (Mumbai region) | Populated · required · entered by a person |
| **Provider documentation *** | https://docs.northwind.example/ai/it-helpdesk-assistant/provider-docs<br>Evidence: provider_documentation https://docs.northwind.example/ai/it-helpdesk-assistant/provider-docs | Populated · required · entered by a person |
| **Audit evidence** | — | Not assessed · recommended |
| **Certifications** | — | Not assessed · recommended |
| **Sub-processors** | — | Not assessed · recommended |
| **Contractual obligations *** | Covered by the bank-wide Asterion agreement. | Populated · required · entered by a person |
| **Incident history** | — | Not assessed · recommended |
| **Change-notification terms *** | 30 days for material changes. | Populated · required · entered by a person |
| **Exit strategy *** | — | Not assessed · required |
| **Concentration risk** | — | Not assessed · recommended |
| **Responsibility split *** | — | Not assessed · required |

## 16. Deployment and change management (0/0 required complete)

| Field | Value | Status |
|---|---|---|
| **Environment** | — | Not assessed · recommended |
| **Deployment date** | — | Not assessed · recommended |
| **Deployment approval** | — | Not assessed · recommended |
| **Change history** | — | Not assessed · recommended |
| **Change-management process** | — | Not assessed · recommended |
| **Decommissioning date** | — | Not assessed · recommended |

## 17. Monitoring and incidents (5/7 required complete)

| Field | Value | Status |
|---|---|---|
| **Monitored dimensions *** | ACCURACY; COST<br>Refs: nist-ai-rmf-1.0:MANAGE-4, eu-ai-act-2024:ART-72 | Populated · required · entered by a person |
| **Thresholds *** | Accuracy below 90% on the weekly sample. | Populated · required · entered by a person |
| **Frequency *** | Weekly | Populated · required · entered by a person |
| **Latest drift status *** | Not yet measured | Populated · required · entered by a person |
| **Incidents *** | {"state":"ZERO_RECORDED","count":0} | Populated · required · entered by a person |
| **Open incidents** | — | Not assessed · recommended |
| **Incident severity and classification** | — | Not assessed · recommended |
| **Affected users** | — | Not assessed · recommended |
| **Root cause** | — | Not assessed · recommended |

DRAFT

| Field | Value | Status |
|---|---|---|
| Corrective action | — | Not assessed · recommended |
| Regulatory notification * | — Refs: eu-ai-act-2024:ART-73, in-dpdp-2023:BREACH_NOTICE | Not assessed · required |
| Suspension criteria * | — | Not assessed · required |
| Remediation status | — | Not assessed · recommended |

### 18. Environmental sustainability (0/0 required complete)

| Field | Value | Status |
|---|---|---|
| Inference energy | — | Not assessed · recommended |
| Carbon emissions | — | Not assessed · recommended |
| Hosting region | — | Not assessed · recommended |
| Efficiency measures | — | Not assessed · recommended |
| Assessment date | — | Not assessed · recommended |

## 3. Risk engine explanation

Score 4 of 23; score tier Low; with floors Low.

| Factor | Value | Points | Note |
|---|---|---|---|
| Impact severity | MINOR | 1/4 | — |
| Likelihood | POSSIBLE | 2/4 | — |
| Exposure (tenants and users reached) | ONE_TENANT | 1/3 | — |
| Autonomy | NONE | 0/3 | — |
| Data sensitivity | LOW | 0/2 | — |
| Decision impact | INFORMATIONAL | 0/4 | — |
| Reversibility | EASILY_REVERSED | 0/3 | — |

## 4. Framework mapping

Where this record supports each framework in the registry. A mapping is not evidence of compliance or certification.

### ISO/IEC 42001

Applies (organisational standard: this card holds supporting evidence only). ISO/IEC 42001: evidence supporting Annex A controls (organisational standard; this card cannot show conformity). Registry revision 1, Unverified requires legal review.

| Control | Title | Status |
|---|---|---|
| A.2 | Annex A group: policies related to AI | Not assessed |
| A.3 | Annex A group: internal organisation, roles and responsibilities | Not assessed |
| A.4 | Annex A group: resources for AI systems | Not assessed |
| A.5 | Annex A group: assessing impacts of AI systems | Not assessed |
| A.6 | Annex A group: AI system life cycle | Not assessed |
| A.7 | Annex A group: data for AI systems | Not assessed |
| A.8 | Annex A group: information for interested parties | Not assessed |
| A.9 | Annex A group: use of AI systems | Not assessed |
| A.10 | Annex A group: third-party and customer relationships | Not assessed |

DRAFT

## NIST AI RMF

Applies. NIST AI RMF: practices mapped to this record (a mapping is not compliance). Registry revision 1, Unverified requires legal review.

| Control | Title | Status |
|---|---|---|
| **GOVERN-2** | Accountability: roles and responsibilities for AI risk are assigned and documented | Partially implemented |
| **MAP-1** | Context: intended purpose, setting and expected users are understood and recorded | Partially implemented |
| **MAP-4** | Third-party components: risks from external models, data and software are mapped | Partially implemented |
| **MAP-5** | Impacts: likelihood and size of impacts on people and organisations are characterised | Not assessed |
| **MEASURE-2** | Evaluation: the system is tested for trustworthiness characteristics, with results recorded | Not assessed |
| **MANAGE-1** | Treatment: risks are prioritised and a response is decided and recorded | Not assessed |
| **MANAGE-4** | Monitoring: risk treatments, incidents and recovery are monitored after deployment | Not assessed |

DRAFT

# 5. Framework sections

One section for each framework enabled for Northwind Finance (sample) in Governance › Setup. Frameworks that are switched off are not shown. Clause results are automated checks of the recorded model fields at the time of generation.

## 5.1 NIST AI RMF 1.0

| Version | Jurisdiction | Model tier (framework) | Met | Gaps (mandatory) | Not assessed |
|---|---|---|---|---|---|
| 1.0.0 | United States | Standard | 12 | 11 (4) | 0 |

| Clause | Requirement | Status | Evidence or gap |
|---|---|---|---|
| **GOVERN 2.1** | Organizational Ownership & Responsibility | Met | Owner team: IT Service Desk |
| **GOVERN 1.2** | Board Policy Alignment | Met | Board policy reference: AI-POL-2026 |
| **GOVERN 1.5** | Audit Cadence Defined | Gap | A recurring audit schedule must be defined for each model. |
| **MAP 1.1** | Purpose & Use-Case Documentation | Met | Use case documented (92 chars) |
| **MAP 1.1** | Deployment Context & Affected Users (optional) | Met | Functional area 'it_operations' affecting segments: Employees |
| **MAP 4.1** | Source & Dependency Provenance (optional) | Met | Provenance: northwind/it/helpdesk-bot.ts |
| **MEASURE 2.7** | Security Red Teaming | Gap | Measure stage requires a red-team review within 365 days for LOW risk (date or attached report). |
| **MEASURE 2.4** | Performance Drift Measurement | Gap | Measure stage requires recent drift evaluation over live inference scores (push windows to /monitoring/scores). |
| **MEASURE 2.3** | Performance Benchmarks Recorded (optional) | Gap | Record baseline performance benchmarks for comparison over time. |
| **GOVERN 3.2** | Human Oversight Controls | Met | Human-in-the-loop oversight enabled |
| **MANAGE 2.4** | Fallback / Business-Continuity Plan (optional) | Gap | High-risk models need a configured business-continuity fallback (activate via /monitoring/bcp). |
| **MANAGE 4.3** | Incident Reporting Hygiene | Met | No overdue incident reports |
| **MANAGE 2.1** | Risk Response Resourcing | Met | Risk response owned by IT Service Desk |
| **MEASURE 3.1** | Risk Tracking Mechanisms | Met | 83 ledger event(s) tracking this model over time |
| **MEASURE 2.10** | Privacy Risk Examined (DPIA / Privacy Review) | Gap | Measure stage requires the system’s privacy risk to be examined and documented (attach a dpia artifact). |
| **MANAGE 1.3** | Risk Responses Planned and Documented (Risk Register) (optional) | Gap | Map and Manage need identified risks with likelihood and impact: record and assess this system's risks. |
| **MAP 3.2** | Context Profile: Costs of Errors, Limitations and Out-of-Scope Uses (optional) | Gap | Map stage needs the system's context documented in structured form: complete the context profile. |
| **GOVERN 6.1** | Third-Party Risk Policies Applied (optional) | Met | No third-party components recorded |
| **MEASURE 2.6** | Safety Re-evaluated After Risk Escalation (optional) | Met | No red-team exercise outstanding |
| **MEASURE 2.13** | Effectiveness of TEVV metrics evaluated (optional) | Gap | Register test sets with acceptance criteria for the system. |
| **MANAGE 4.3** | Incidents communicated to relevant AI actors (optional) | Met | — |
| **MEASURE 2.12** | Environmental impact and sustainability assessed (optional) | Gap | Record the system's energy use (kWh) and CO2e. |
| **MEASURE 2.7** | AI system security and resilience evaluated (attack categories) (optional) | Gap | Approve an AI threat model for the system (state: NONE). |

DRAFT

## 5.2 ISO/IEC 42001:2023

| Version | Jurisdiction | Model tier (framework) | Met | Gaps (mandatory) | Not assessed |
|---|---|---|---|---|---|
| 2023.1.0 | International | Standard | 6 | 12 (0) | 1 |

| Clause | Requirement | Status | Evidence or gap |
|---|---|---|---|
| **6.1.2** | AI System Risk Assessment | Met | System evaluated risk tier: LOW |
| **8.4** | AI System Impact Assessment (6.1.4 / 8.4) (optional) | Gap | Clauses 6.1.4 and 8.4 require a signed AI system impact assessment: complete and finalise one for this model. |
| **A.7.5** | Data & Codebase Provenance | Met | Source/data references tracked: northwind/it/helpdesk-bot.ts |
| **9.1** | Performance Monitoring & Measurement (optional) | Gap | Clause 9.1 requires active performance monitoring of deployed AI systems. |
| **9.2** | Internal Audit Cadence (optional) | Gap | Clause 9.2 requires internal audits at planned intervals; schedule the next audit. |
| **10.2** | Nonconformity & Corrective Action (optional) | Met | No overdue incident reports |
| **5.3** | Roles, Responsibilities & Accountability | Met | Accountable team assigned: IT Service Desk |
| **A.6.2.7** | AI System Technical Documentation (optional) | Gap | Annex A.6.2.7 recommends technical documentation for relevant parties (attach a model_card artifact). |
| **6.1.3** | AI Risk Treatment (Risk Register) (optional) | Gap | Clauses 6.1.2–6.1.3 and 8.2–8.3 need an assessed risk register: record and assess this system's risks. |
| **A.6.2.2** | AI System Requirements and Specification (optional) | Gap | Annex A.6.2.2 needs the system's requirements and intended use specified: complete the context profile. |
| **A.8.2** | System Documentation and Information for Users (optional) | Gap | Annex A.8.2 needs information for users, including limitations: complete the context profile. |
| **A.10.3** | Suppliers (optional) | Met | No third-party components recorded |
| **A.7.2** | Data for Development and Enhancement (optional) | Gap | A.7 needs the data used by the system defined and documented: link its datasets. |
| **A.6.2.4** | Verification and Validation (optional) | Gap | A.6.2.4 needs verification and validation of the system recorded. |
| **A.6.2.4** | Verification and Validation (test sets) (optional) | Gap | Register test sets with acceptance criteria for the system. |
| **A.8.4** | Communication of Incidents (optional) | Met | — |
| **A.5.5** | Societal Impacts (environmental) (optional) | Gap | Record the system's energy use (kWh) and CO2e. |
| **A.6.2.6** | Operation and Monitoring (AI security) (optional) | Gap | Approve an AI threat model for the system (state: NONE). |
| **A.10.4** | Customers (optional) | Not assessed | — |

# 6. Approval history

No approvals recorded.

DRAFT

# 7. Document control

| Attribute | Value |
|---|---|
| **Status** | DRAFT (live) |
| **Workflow state** | Evidence collection |
| **Organisation** | Northwind Finance (sample) · Lending and operations AI |
| **Record ID** | aebae446-540b-4f83-a27f-62eae9aa0887 |
| **Content hash (SHA-256)** | ea9e4347ee69de2f2f3242f7fd5c5cfbd57ba50513b7363a4049131abd5fa482 |
| **Generated** | 2026-11-27T15:17:15.757Z by asha.menon@northwind.example |
| **Catalogue version** | 2.0.0 |

*This model card documents the system. Framework mappings are not evidence of compliance or certification.*

A model card is a documentation and transparency artefact. It is not, by itself, proof of regulatory compliance, certification, or conformity of an organisation's AI management system.

DRAFT