SENTRAI | GOVERNANCE

Northwind Finance (sample)
ORGANISATION

AI MODEL GOVERNANCE SUMMARY

BRANCH FAQ CHATBOT | 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.

0 / 77	0 / 23	0 / 16	0 / 2
DOCUMENTATION	EVIDENCE	CONTROL EFFECTIVENESS	REGULATORY ASSESSMENT
0% complete	0% complete	0% complete	0% complete

RECORD AT A GLANCE

Attribute	Recorded value	Attribute	Recorded value
Model / record name	Branch FAQ Chatbot	**Record type**	—
Component lifecycle	—	**Model lifecycle**	Design
Deployment status	Not deployed	**Operational status**	Active
Inherent risk tier	—	**Residual risk tier**	— (Not assessed)
EU AI Act classification	—	**Korea high-impact**	—
Workflow state	Draft	**Next review**	— (Not scheduled)

FRAMEWORK STATUS

Framework (enabled for the organisation)	Automated clause check
NIST AI RMF 1.0	6 open gaps
ISO/IEC 42001:2023	1 open gap

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

RISK POSTURE

Inherent —, residual — (Not assessed). Risk engine: score 23 of 23, score tier Critical, Critical with floors. A person sets the tiers; the engine explains and checks them.

KEY GOVERNANCE ATTENTION POINTS

Priority	Finding	Why it matters
High	81 items block a final card	Record type is required. Model or record identifier is required. Name: a suggested or imported value is waiting for a person to accept or edit it. See section 1.
High	Risk engine finding	A person must set (or accept) the inherent risk tier.
High	Risk engine finding	A person must set (or accept) the residual risk tier.
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: 6 open gaps	Mandatory clauses not yet met: GOVERN 2.1, GOVERN 1.5, MEASURE 2.7, MEASURE 2.4, ….

DRAFT

Priority	Finding	Why it matters
Medium	ISO/IEC 42001:2023: 1 open gap	Mandatory clauses not yet met: 5.3.
Low	No audit scheduled	Next audit date is not set.

DRAFT

1. Finalisation gate

#	Blocking a final card
1	Record type is required.
2	Model or record identifier is required.
3	Name: a suggested or imported value is waiting for a person to accept or edit it.
4	Provider is required.
5	Model category: a suggested or imported value is waiting for a person to accept or edit it.
6	Version is required.
7	Hosting type is required.
8	Hosting region is required.
9	Component lifecycle is required.
10	Model lifecycle: a suggested or imported value is waiting for a person to accept or edit it.
11	Deployment status: a suggested or imported value is waiting for a person to accept or edit it.
12	Operational status: a suggested or imported value is waiting for a person to accept or edit it.
13	Accountable owner is required.
14	Technical owner is required.
15	Next review date is required.
16	Business justification is required.
17	Purpose: a suggested or imported value is waiting for a person to accept or edit it.
18	Intended use is required.
19	Permitted use is required.
20	Intended users is required.
21	Affected segments is required.
22	Capabilities is required.
23	Geographic scope is required.
24	Decision impact is required.
25	Out-of-scope uses is required.
26	Prohibited uses is required.
27	Known limitations is required.
28	Inference configuration is required (required because R_GENERATIVE fired).
29	Prompts and system prompt is required (required because R_GENERATIVE fired).
30	Data classification is required.
31	Personal data processed: a suggested or imported value is waiting for a person to accept or edit it.
32	Third-party provider used is required.
33	Models this configuration can select is required.
34	Impact severity is required.
35	Likelihood is required.
36	Users reached is required.
37	Tenant reach is required.
38	Autonomy level is required.
39	Data sensitivity is required.

DRAFT

#	Blocking a final card
40	Reversibility is required.
41	Inherent risk tier: a suggested or imported value is waiting for a person to accept or edit it.
42	Residual risk tier is required.
43	Risk rationale is required.
44	User disclosure is required (required because R_GENERATIVE fired).
45	AI-content labelling is required (required because R_GENERATIVE fired).
46	Oversight model is required.
47	Responsible role is required.
48	Who may change the setting is required.
49	Four-eyes approval is required.
50	Change audit trail is required.
51	Rollback capability is required.
52	Evaluation records is required.
53	Hallucination evaluation is required (required because R_GENERATIVE fired).
54	Harmful-output evaluation is required (required because R_GENERATIVE fired).
55	Prompt-injection testing is required (required because R_GENERATIVE fired).
56	Regression evaluation before a default switch is required.
57	Next revalidation date is required.
58	Authentication is required.
59	Authorisation is required.
60	Prompt-injection protection is required (required because R_GENERATIVE fired).
61	Incident response is required.
62	Moderation is required (required because R_GENERATIVE fired).
63	Output filtering is required (required because R_GENERATIVE fired).
64	Guardrails is required (required because R_GENERATIVE fired).
65	Safety testing is required (required because R_GENERATIVE fired).
66	Jurisdictions is required.
67	Applicable laws is required.
68	Applicability rationale is required.
69	Compliance owner is required.
70	Legal review date is required.
71	Default model is required.
72	Tenant overrides (BYOK) is required.
73	Model selection rules is required.
74	Tenants on the default is required.
75	Provider and API endpoints is required.
76	Configuration version is required.
77	Incidents is required.
78	Model lifecycle is DESIGN but operational status is ACTIVE. A record that has not been deployed cannot be operating.
79	Deployment status is NOT_DEPLOYED but operational status is ACTIVE.
80	A person must set the inherent and residual risk tiers.

DRAFT

#	Blocking a final card
81	Name an accountable owner: an active person in the organisation.

Warnings

#	Warning
1	Developer is recommended.
2	Model family is recommended.
3	Algorithm or method is recommended.
4	Architecture is recommended.
5	Release date is recommended.
6	Fingerprint or hash is recommended.
7	Licence is recommended.
8	Access type is recommended.
9	Deployment identifier is recommended.
10	Business owner is recommended.
11	Registration date is recommended.
12	Last updated is recommended.
13	Observed use is recommended.
14	Target population is recommended.
15	Tasks is recommended.
16	Languages is recommended.
17	Operating environment is recommended.
18	Expected outputs is recommended.
19	Foreseeable misuse is recommended.
20	Expected performance is recommended.
21	Usage restrictions is recommended.
22	Architecture description is recommended.
23	Parameter count is recommended.
24	Input modalities is recommended.
25	Output modalities is recommended.
26	Context window is recommended.
27	Token limits is recommended.
28	Formats is recommended.
29	Fine-tuning is recommended.
30	Hardware is recommended.
31	Runtime is recommended.
32	API integration is recommended.
33	Hosting and deployment architecture is recommended.
34	Serving technology is recommended.
35	Training data sources is recommended.
36	Training data period is recommended.
37	Data ownership is recommended.
38	Data lineage is recommended.

DRAFT

#	Warning
39	Data quality is recommended.
40	Representativeness is recommended.
41	Personal-data categories is recommended.
42	Personal-data flows is recommended.
43	Sensitive categories is recommended.
44	Retention is recommended.
45	Data residency is recommended.
46	Cross-border transfers is recommended.
47	Transfer destination countries is recommended.
48	Data minimisation is recommended.
49	Lawful basis or consent is recommended.
50	Privacy impact assessment is recommended.
51	Copyright and IP is recommended.
52	Data-poisoning assessment is recommended.
53	Synthetic data is recommended.
54	Deletion mechanism is recommended.
55	Provider may train on inputs is recommended.
56	Knowledge-base sources is recommended.
57	Ingestion is recommended.
58	Chunking is recommended.
59	Embedding model is recommended.
60	Vector database is recommended.
61	Retrieval method is recommended.
62	Retrieval evaluation is recommended.
63	Access-control enforcement is recommended.
64	Refresh frequency is recommended.
65	Indirect prompt-injection controls is recommended.
66	Foundation, parent or fine-tuned-from model is recommended.
67	Embedding dependencies is recommended.
68	External APIs is recommended.
69	Third-party services is recommended.
70	Sub-processors is recommended.
71	Open-source libraries and versions is recommended.
72	SBOM or AI-BOM reference is recommended.
73	Provider documentation is recommended.
74	SLAs is recommended.
75	Data processing agreements is recommended.
76	Provider certifications is recommended.
77	Third-party risk assessment is recommended.
78	Model update policy is recommended.
79	End-of-life and deprecation policy is recommended.

DRAFT

#	Warning
80	Organisations reached is recommended.
81	Tenants reached is recommended.
82	Existing controls is recommended.
83	Risk acceptance is recommended.
84	Impact assessment is recommended.
85	Human and fundamental rights impact is recommended.
86	Privacy impact is recommended.
87	Societal impact is recommended.
88	Financial impact is recommended.
89	Employment impact is recommended.
90	Health and safety impact is recommended.
91	Discrimination impact is recommended.
92	Vulnerable groups is recommended.
93	Environmental impact is recommended.
94	Stakeholder consultation is recommended.
95	Mitigations is recommended.
96	Impact assessment approval is recommended.
97	Fundamental rights impact assessment (EU AI Act Art. 27) is recommended.
98	ISO/IEC 42005 alignment is recommended.
99	Methodology is recommended.
100	Relevant groups is recommended.
101	Representativeness is recommended.
102	Metrics is recommended.
103	Disparate impact is recommended.
104	Results is recommended.
105	Mitigation is recommended.
106	Residual bias is recommended.
107	Review date is recommended.
108	Independent evidence is recommended.
109	Explainability methodology is recommended.
110	Interpretability is recommended.
111	Explainability limitations is recommended.
112	Confidence information is recommended.
113	Traceability is recommended.
114	Notification mechanisms is recommended.
115	Appeal and correction is recommended.
116	Review frequency is recommended.
117	Intervention and override is recommended.
118	Escalation is recommended.
119	Reviewer competency is recommended.
120	AI literacy training is recommended.

DRAFT

#	Warning
121	Final decision accountability is recommended.
122	Automation boundaries is recommended.
123	Tenant notification is recommended.
124	Tools is recommended.
125	Permissions is recommended.
126	Execution privileges is recommended.
127	Autonomous actions is recommended.
128	Approval gates is recommended.
129	Prohibited actions is recommended.
130	Transaction limits is recommended.
131	Emergency stop is recommended.
132	Cross-tenant leakage testing is recommended.
133	Access control is recommended.
134	API-key and BYOK key management is recommended.
135	Encryption at rest is recommended.
136	Encryption in transit is recommended.
137	Model-extraction protection is recommended.
138	Adversarial resistance is recommended.
139	Poisoning protection is recommended.
140	Supply-chain security is recommended.
141	Vulnerability management is recommended.
142	Rate limiting is recommended.
143	Abuse prevention is recommended.
144	Disaster recovery is recommended.
145	Rollback is recommended.
146	Business continuity is recommended.
147	Security assessment evidence is recommended.
148	Safety policies is recommended.
149	Harm-prevention controls is recommended.
150	Misuse prevention is recommended.
151	Dangerous-capability assessment is recommended.
152	Robustness safeguards is recommended.
153	Safety incidents is recommended.
154	Residual safety risk is recommended.
155	EU AI Act role is recommended.
156	EU AI Act classification is recommended.
157	EU AI Act Art. 27 criteria met is recommended.
158	Korea: high-impact AI is recommended.
159	India: acts as an intermediary is recommended.
160	China: generative AI service filing is recommended.
161	US: consequential decision is recommended.

DRAFT

#	Warning
162	Other regulatory classifications is recommended.
163	Exemptions is recommended.
164	Next regulatory review is recommended.
165	Provider is recommended.
166	Provider location is recommended.
167	Provider documentation is recommended.
168	Audit evidence is recommended.
169	Certifications is recommended.
170	Sub-processors is recommended.
171	Data processing agreement is recommended.
172	Contractual obligations is recommended.
173	Incident history is recommended.
174	Change-notification terms is recommended.
175	Exit strategy is recommended.
176	Concentration risk is recommended.
177	Responsibility split is recommended.
178	Environment is recommended.
179	Deployment date is recommended.
180	Deployment approval is recommended.
181	Change history is recommended.
182	Change-management process is recommended.
183	Decommissioning date is recommended.
184	Users on the default is recommended.
185	Downstream dependencies is recommended.
186	Open incidents is recommended.
187	Incident severity and classification is recommended.
188	Affected users is recommended.
189	Root cause is recommended.
190	Corrective action is recommended.
191	Regulatory notification is recommended.
192	Grievance channel is recommended.
193	Grievances is recommended.
194	Grievance resolution is recommended.
195	Remediation status is recommended.
196	Training energy is recommended.
197	Inference energy is recommended.
198	Carbon emissions is recommended.
199	Water use is recommended.
200	Hardware and GPU utilisation is recommended.
201	Hosting region is recommended.
202	Efficiency measures is recommended.

DRAFT

#	Warning
203	Assessment date is recommended.
204	Counted as the worst plausible value until assessed: impact_severity, likelihood, tenant_reach, users_reached, autonomy_level, data_sensitivity, decision_impact, reversibility.

2. Card sections

1. Identity and metadata (0/15 required complete)

Field	Value	Status
Record type *	—	Not assessed · required
Model or record identifier *	—	Not assessed · required
Name *	Branch FAQ Chatbot	Populated · required · imported from v1 · awaiting acceptance
Provider *	—	Not assessed · required
Developer	unknown	Populated · recommended · imported from v1 · awaiting acceptance
Model family	—	Not assessed · recommended
Model category *	GENERATIVE	Populated · required · CODE SCAN, not accepted · awaiting acceptance
Algorithm or method	—	Not assessed · recommended
Architecture	—	Not assessed · recommended
Version *	—	Not assessed · required
Release date	—	Not assessed · recommended
Fingerprint or hash	—	Not assessed · recommended
Licence	—	Not assessed · recommended
Access type	—	Not assessed · recommended
Hosting type *	—	Not assessed · required
Hosting region *	—	Not assessed · required
Deployment identifier	—	Not assessed · recommended
Component lifecycle *	—	Not assessed · required
Model lifecycle *	DESIGN	Populated · required · imported from v1 · awaiting acceptance
Deployment status *	NOT_DEPLOYED	Populated · required · imported from v1 · awaiting acceptance
Operational status *	ACTIVE	Populated · required · imported from v1 · awaiting acceptance
Accountable owner *	— Refs: nist-ai-rmf-1.0:GOVERN-2, iso-42001-2023:A.3	Not assessed · required
Technical owner *	—	Not assessed · required
Business owner	—	Not assessed · recommended
Registration date	2026-11-27	Populated · recommended · imported from v1 · awaiting acceptance
Last updated	—	Not assessed · recommended
Next review date *	—	Not assessed · required

2. Intended use and scope (0/12 required complete)

DRAFT

Field	Value	Status
Business justification *	—	Not assessed · required
Purpose *	Here is a concise summary of the component: The chatbot answers customer questions about branch hours and products on the public website. It likely handles sensitive information. Note: This summary is based on the provided code excerpt. Refs: nist-ai-rmf-1.0:MAP-1, iso-42001-2023:A.6	Populated · required · imported from v1 · awaiting acceptance
Intended use *	—	Not assessed · required
Permitted use *	—	Not assessed · required
Observed use	—	Not assessed · recommended
Intended users *	—	Not assessed · required
Affected segments *	— Refs: nist-ai-rmf-1.0:MAP-5	Not assessed · required
Target population	—	Not assessed · recommended
Tasks	—	Not assessed · recommended
Capabilities *	—	Not assessed · required
Languages	—	Not assessed · recommended
Geographic scope *	—	Not assessed · required
Operating environment	—	Not assessed · recommended
Expected outputs	—	Not assessed · recommended
Decision impact *	—	Not assessed · required
Out-of-scope uses *	—	Not assessed · required
Prohibited uses *	—	Not assessed · required
Foreseeable misuse	— Refs: eu-ai-act-2024:ART-9	Not assessed · recommended
Known limitations *	—	Not assessed · required
Expected performance	—	Not assessed · recommended
Usage restrictions	—	Not assessed · recommended

3. Technical specifications (0/2 required complete)

Field	Value	Status
Architecture description	—	Not assessed · recommended
Parameter count	—	Not assessed · recommended
Input modalities	—	Not assessed · recommended
Output modalities	—	Not assessed · recommended
Context window	—	Not assessed · recommended
Token limits	—	Not assessed · recommended
Formats	—	Not assessed · recommended
Inference configuration *	—	Not assessed · required
Prompts and system prompt *	—	Not assessed · required
Fine-tuning	—	Not assessed · recommended
Hardware	—	Not assessed · recommended
Runtime	—	Not assessed · recommended

DRAFT

Field	Value	Status
API integration	—	Not assessed · recommended
Hosting and deployment architecture	—	Not assessed · recommended
Serving technology	—	Not assessed · recommended

4. Data governance and provenance (0/2 required complete)

Field	Value	Status
Training data sources	— 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 *	—	Not assessed · required
Data quality	—	Not assessed · recommended
Representativeness	—	Not assessed · recommended
Personal data processed *	UNKNOWN Refs: gdpr:ART-6, in-dpdp-2023:NOTICE	Unknown · required · imported from v1 · awaiting acceptance
Personal-data categories	—	Not assessed · recommended
Personal-data flows	—	Not assessed · recommended
Sensitive categories	— Refs: gdpr:ART-9	Not assessed · recommended
Retention	— Refs: in-dpdp-2023:RETENTION	Not assessed · recommended
Data residency	—	Not assessed · recommended
Cross-border transfers	— Refs: gdpr:ART-44, in-dpdp-2023:CROSS_BORDER	Not assessed · recommended
Transfer destination countries	—	Not assessed · recommended
Data minimisation	—	Not assessed · recommended
Lawful basis or consent	— Refs: gdpr:ART-6, in-dpdp-2023:CONSENT	Not assessed · recommended
Privacy impact assessment	— Refs: gdpr:ART-35	Not assessed · recommended
Copyright and IP	— Refs: nist-ai-600-1:INTELLECTUAL_PROPERTY	Not assessed · recommended
Data-poisoning assessment	—	Not assessed · recommended
Synthetic data	—	Not assessed · recommended
Deletion mechanism	—	Not assessed · recommended
Provider may train on inputs	— Refs: nist-ai-600-1:DATA_PRIVACY	Not assessed · recommended
Knowledge-base sources	—	Not assessed · recommended
Ingestion	—	Not assessed · recommended
Chunking	—	Not assessed · recommended
Embedding model	—	Not assessed · recommended
Vector database	—	Not assessed · recommended
Retrieval method	—	Not assessed · recommended

DRAFT

Field	Value	Status
Retrieval evaluation	—	Not assessed · recommended
Access-control enforcement	—	Not assessed · recommended
Refresh frequency	—	Not assessed · recommended
Indirect prompt-injection controls	— Refs: nist-ai-600-1:INFORMATION_SECURITY	Not assessed · recommended

5. Dependencies and supply chain (0/2 required complete)

Field	Value	Status
Third-party provider used *	—	Not assessed · required
Foundation, parent or fine-tuned-from model	—	Not assessed · recommended
Models this configuration can select *	—	Not assessed · required
Embedding dependencies	—	Not assessed · recommended
External APIs	—	Not assessed · recommended
Third-party services	—	Not assessed · recommended
Sub-processors	—	Not assessed · recommended
Open-source libraries and versions	—	Not assessed · recommended
SBOM or AI-BOM reference	—	Not assessed · recommended
Provider documentation	—	Not assessed · recommended
SLAs	—	Not assessed · recommended
Data processing agreements	—	Not assessed · recommended
Provider certifications	—	Not assessed · recommended
Third-party risk assessment	— Refs: iso-42001-2023:A.10, nist-ai-600-1:VALUE_CHAIN	Not assessed · recommended
Model update policy	—	Not assessed · recommended
End-of-life and deprecation policy	—	Not assessed · recommended

6. Risk classification (0/10 required complete)

Field	Value	Status
Impact severity *	—	Not assessed · required
Likelihood *	—	Not assessed · required
Users reached *	—	Not assessed · required
Organisations reached	—	Not assessed · recommended
Tenant reach *	—	Not assessed · required
Tenants reached	—	Not assessed · recommended
Autonomy level *	—	Not assessed · required
Data sensitivity *	—	Not assessed · required
Reversibility *	—	Not assessed · required
Existing controls	—	Not assessed · recommended
Inherent risk tier *	LOW Refs: nist-ai-rmf-1.0:MAP-4, iso-23894-2023:RISK_PROCESS	Populated · required · imported from v1 · awaiting acceptance

DRAFT

Field	Value	Status
Residual risk tier *	— Refs: nist-ai-rmf-1.0:MANAGE-1	Not assessed · required
Risk rationale *	—	Not assessed · required
Risk acceptance	—	Not assessed · recommended

7. Impact assessment (0/0 required complete)

Field	Value	Status
Impact assessment	— Refs: iso-42001-2023:A.5, iso-42005-2025:IMPACT_ASSESSMENT	Not assessed · recommended
Human and fundamental rights impact	—	Not assessed · recommended
Privacy impact	—	Not assessed · recommended
Societal impact	—	Not assessed · recommended
Financial impact	—	Not assessed · recommended
Employment impact	—	Not assessed · recommended
Health and safety impact	—	Not assessed · recommended
Discrimination impact	—	Not assessed · recommended
Vulnerable groups	—	Not assessed · recommended
Environmental impact	—	Not assessed · recommended
Stakeholder consultation	—	Not assessed · recommended
Mitigations	—	Not assessed · recommended
Impact assessment approval	—	Not assessed · recommended
Fundamental rights impact assessment (EU AI Act Art. 27)	— Refs: eu-ai-act-2024:ART-27	Not assessed · recommended
ISO/IEC 42005 alignment	— Refs: iso-42005-2025:IMPACT_ASSESSMENT	Not assessed · recommended

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

Field	Value	Status
Methodology	— Refs: nist-ai-rmf-1.0:MEASURE-2, nist-ai-600-1:HARMFUL_BIAS	Not assessed · recommended
Relevant groups	—	Not assessed · recommended
Representativeness	—	Not assessed · recommended
Metrics	—	Not assessed · recommended
Disparate impact	—	Not assessed · recommended
Results	—	Not assessed · recommended
Mitigation	—	Not assessed · recommended
Residual bias	—	Not assessed · recommended
Review date	—	Not assessed · recommended
Independent evidence	—	Not assessed · recommended

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

DRAFT

Field	Value	Status
Explainability methodology	— Refs: nist-ai-rmf-1.0:MEASURE-2	Not assessed · recommended
Interpretability	—	Not assessed · recommended
Explainability limitations	—	Not assessed · recommended
Confidence information	—	Not assessed · recommended
Traceability	— Refs: eu-ai-act-2024:ART-12	Not assessed · recommended
User disclosure *	— Refs: eu-ai-act-2024:ART-50, kr-ai-basic-act:USER_NOTICE	Not assessed · required
AI-content labelling *	— 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
Appeal and correction	—	Not assessed · recommended

10. Human oversight and accountability (0/6 required complete)

Field	Value	Status
Oversight model *	— Refs: eu-ai-act-2024:ART-14, nist-ai-600-1:HUMAN_AI_CONFIGURATION	Not assessed · required
Responsible role *	—	Not assessed · required
Review frequency	—	Not assessed · recommended
Intervention and override	—	Not assessed · recommended
Escalation	—	Not assessed · recommended
Reviewer competency	—	Not assessed · recommended
AI literacy training	— Refs: eu-ai-act-2024:ART-4	Not assessed · recommended
Final decision accountability	—	Not assessed · recommended
Automation boundaries	—	Not assessed · recommended
Who may change the setting *	—	Not assessed · required
Four-eyes approval *	—	Not assessed · required
Change audit trail *	—	Not assessed · required
Tenant notification	—	Not assessed · recommended
Rollback capability *	—	Not assessed · required
Tools	—	Not assessed · recommended
Permissions	—	Not assessed · recommended
Execution privileges	—	Not assessed · recommended
Autonomous actions	—	Not assessed · recommended
Approval gates	—	Not assessed · recommended
Prohibited actions	—	Not assessed · recommended
Transaction limits	—	Not assessed · recommended
Emergency stop	—	Not assessed · recommended

11. Evaluation and validation (0/6 required complete)

DRAFT

Field	Value	Status
Evaluation records *	— Refs: nist-ai-rmf-1.0:MEASURE-2, eu-ai-act-2024:ART-15	Not assessed · required
Hallucination evaluation *	— Refs: nist-ai-600-1:CONFABULATION	Not assessed · required
Harmful-output evaluation *	— Refs: nist-ai-600-1:DANGEROUS_CONTENT	Not assessed · required
Prompt-injection testing *	— Refs: nist-ai-600-1:INFORMATION_SECURITY	Not assessed · required
Cross-tenant leakage testing	—	Not assessed · recommended
Regression evaluation before a default switch *	—	Not assessed · required
Next revalidation date *	—	Not assessed · required

12. Security and resilience (0/4 required complete)

Field	Value	Status
Authentication *	—	Not assessed · required
Authorisation *	—	Not assessed · required
Access control	—	Not assessed · recommended
API-key and BYOK key management	—	Not assessed · recommended
Encryption at rest	—	Not assessed · recommended
Encryption in transit	—	Not assessed · recommended
Prompt-injection protection *	— Refs: nist-ai-600-1:INFORMATION_SECURITY	Not assessed · required
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	— 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

DRAFT

Field	Value	Status
Misuse prevention	—	Not assessed · recommended
Dangerous-capability assessment	— 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 (0/5 required complete)

Field	Value	Status
Jurisdictions *	—	Not assessed · required
EU AI Act role	— Refs: eu-ai-act-2024:ART-3	Not assessed · recommended
EU AI Act classification	— Refs: eu-ai-act-2024:ART-5, eu-ai-act-2024:ART-6	Not assessed · recommended
EU AI Act Art. 27 criteria met	—	Not assessed · recommended
Korea: high-impact AI	— Refs: kr-ai-basic-act:HIGH_IMPACT	Not assessed · recommended
India: acts as an intermediary	—	Not assessed · recommended
China: generative AI service filing	— Refs: cn-genai-filing:FILING	Not assessed · recommended
US: consequential decision	— Refs: us-co-sb26-189:CONSEQUENTIAL_DECISION	Not assessed · recommended
Other regulatory classifications	—	Not assessed · recommended
Applicable laws *	—	Not assessed · required
Applicability rationale *	—	Not assessed · required
Exemptions	—	Not assessed · recommended
Compliance owner *	—	Not assessed · required
Legal review date *	—	Not assessed · required
Next regulatory review	—	Not assessed · recommended

15. Third-party governance (0/0 required complete)

Field	Value	Status
Provider	— Refs: iso-42001-2023:A.10	Not assessed · recommended
Provider location	—	Not assessed · recommended
Provider documentation	—	Not assessed · recommended
Audit evidence	—	Not assessed · recommended
Certifications	—	Not assessed · recommended
Sub-processors	—	Not assessed · recommended
Data processing agreement	—	Not assessed · recommended
Contractual obligations	—	Not assessed · recommended
Incident history	—	Not assessed · recommended
Change-notification terms	—	Not assessed · recommended

DRAFT

Field	Value	Status
Exit strategy	—	Not assessed · recommended
Concentration risk	—	Not assessed · recommended
Responsibility split	—	Not assessed · recommended

16. Deployment and change management (0/6 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
Default model *	—	Not assessed · required
Tenant overrides (BYOK) *	—	Not assessed · required
Model selection rules *	—	Not assessed · required
Tenants on the default *	—	Not assessed · required
Users on the default	—	Not assessed · recommended
Provider and API endpoints *	—	Not assessed · required
Configuration version *	—	Not assessed · required
Downstream dependencies	—	Not assessed · recommended

17. Monitoring and incidents (0/1 required complete)

Field	Value	Status
Incidents *	—	Not assessed · required
Open incidents	—	Not assessed · recommended
Incident severity and classification	—	Not assessed · recommended
Affected users	—	Not assessed · recommended
Root cause	—	Not assessed · recommended
Corrective action	—	Not assessed · recommended
Regulatory notification	— Refs: eu-ai-act-2024:ART-73, in-dpdp-2023:BREACH_NOTICE	Not assessed · recommended
Grievance channel	—	Not assessed · recommended
Grievances	—	Not assessed · recommended
Grievance resolution	—	Not assessed · recommended
Remediation status	—	Not assessed · recommended

18. Environmental sustainability (0/0 required complete)

Field	Value	Status
Training energy	— Refs: nist-ai-600-1:ENVIRONMENTAL	Not assessed · recommended
Inference energy	—	Not assessed · recommended
Carbon emissions	—	Not assessed · recommended

DRAFT

Field	Value	Status
Water use	—	Not assessed · recommended
Hardware and GPU utilisation	—	Not assessed · recommended
Hosting region	—	Not assessed · recommended
Efficiency measures	—	Not assessed · recommended
Assessment date	—	Not assessed · recommended

3. Risk engine explanation

Score 23 of 23; score tier Critical; with floors Critical.

Factor	Value	Points	Note
Impact severity	—	4/4	Not assessed: counted as SEVERE (the worst plausible value) until a person records it.
Likelihood	—	4/4	Not assessed: counted as ALMOST_CERTAIN (the worst plausible value) until a person records it.
Exposure (tenants and users reached)	—	3/3	Not assessed: counted as ALL_TENANTS (the worst plausible value) until a person records it.
Autonomy	—	3/3	Not assessed: counted as ACTS_AUTONOMOUSLY (the worst plausible value) until a person records it.
Data sensitivity	—	2/2	Not assessed: counted as HIGH (the worst plausible value) until a person records it.
Decision impact	—	4/4	Not assessed: counted as SAFETY_CRITICAL (the worst plausible value) until a person records it.
Reversibility	—	3/3	Not assessed: counted as IRREVERSIBLE (the worst plausible value) until a person records it.
Floor HIGH	—	—	Impact severity is not assessed, so it counts as SEVERE: at least HIGH.
Blocking finding	—	—	A person must set (or accept) the inherent risk tier.
Blocking finding	—	—	A person must set (or accept) the residual risk tier.
Note	rule 2	—	Counted as the worst plausible value until assessed: impact_severity, likelihood, tenant_reach, users_reached, autonomy_level, data_sensitivity, decision_impact, reversibility.

4. Framework mapping

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

China AI labelling

Applicability not yet determined. China AI labelling: obligations mapped to this record (a mapping is not compliance). Registry revision 1, Unverified requires legal review.

Control	Title	Status
LABEL	AI-generated content carries explicit and implicit labels	Not assessed (date to be confirmed)

China GenAI filing

Applicability not yet determined. China GenAI filing: obligations mapped to this record (a mapping is not compliance). Registry revision 1, Unverified requires legal review.

Control	Title	Status
FILING	A generative AI service offered to the public is filed with the regulator	Not assessed (date to be confirmed)

EU AI Act

Applicability not yet determined. EU AI Act: obligations mapped to this record (a mapping is not compliance). Registry revision 1, Unverified requires legal review.

Control	Title	Status
ART-3	Determine the organisation's role for this system (provider, deployer, importer, distributor)	Not assessed
ART-4	Staff dealing with the system have a sufficient level of AI literacy	Not assessed
ART-5	The use is not a prohibited practice	Not assessed
ART-6	High-risk classification is decided and recorded	Not assessed
ART-9	High-risk: a risk management system runs across the life cycle	Not assessed (date to be confirmed)
ART-10	High-risk: training, validation and test data meet quality and governance requirements	Not assessed (date to be confirmed)
ART-12	High-risk: events are logged automatically over the system's lifetime	Not assessed (date to be confirmed)
ART-14	High-risk: designed for effective human oversight	Not assessed (date to be confirmed)
ART-15	High-risk: appropriate accuracy, robustness and cybersecurity	Not assessed (date to be confirmed)
ART-27	Certain deployers of high-risk systems assess the impact on fundamental rights	Not assessed (date to be confirmed)
ART-50	Transparency: people are told they interact with AI, and synthetic content is marked	Not assessed
ART-53	GPAI model providers keep technical documentation and a copyright policy, and inform downstream providers	Not assessed (date to be confirmed)
ART-55	GPAI models with systemic risk: evaluation, incident reporting and cybersecurity duties	Not assessed (date to be confirmed)
ART-72	High-risk: providers run post-market monitoring	Not assessed (date to be confirmed)
ART-73	High-risk: serious incidents are reported to the authorities	Not assessed (date to be confirmed)

GDPR

Applicability not yet determined. GDPR: obligations mapped to this record (a mapping is not compliance). Registry revision 1, Unverified requires legal review.

Control	Title	Status
ART-6	A lawful basis is identified for the processing	Not assessed
ART-9	Special categories of data are processed only under a permitted condition	Not assessed
ART-35	A data protection impact assessment is done where the processing is likely high risk	Not assessed
ART-44	Transfers outside the EEA rely on a permitted mechanism	Not assessed

DPDP Act

Applicability not yet determined. DPDP Act: obligations mapped to this record (a mapping is not compliance). Registry revision 1, Unverified requires legal review.

Control	Title	Status
NOTICE	Data principals get a notice of what is processed and why	Not assessed (date to be confirmed)
CONSENT	Processing rests on valid consent or a permitted legitimate use	Not assessed (date to be confirmed)
RETENTION	Personal data is erased when the purpose is served	Not assessed (date to be confirmed)
CROSS_BORDER	Transfers abroad respect any government restriction	Not assessed (date to be confirmed)
BREACH_NOTICE	Personal data breaches are notified to the Board and the people affected	Not assessed (date to be confirmed)

DPDP Rules

Applicability not yet determined. DPDP Rules: obligations mapped to this record (a mapping is not compliance). Registry revision 1, Unverified requires legal review.

IT Rules 2026

Applicability not yet determined. IT Rules 2026: obligations mapped to this record (a mapping is not compliance). Registry revision 1, Unverified requires legal review.

DRAFT

Control	Title	Status
SYNTHETIC_LABEL	Synthetically generated content is labelled	Not assessed (date to be confirmed)
TAKEDOWN	Takedown requests are acted on within the required time	Not assessed (date to be confirmed)

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

Korea AI Basic Act

Applicability not yet determined. Korea AI Basic Act: obligations mapped to this record (a mapping is not compliance). Registry revision 1, Unverified requires legal review.

Control	Title	Status
HIGH_IMPACT	Whether the system is high-impact AI is decided and recorded	Not assessed
USER_NOTICE	Users are told in advance when generative or high-impact AI is used	Not assessed
OUTPUT_LABEL	Generative outputs are labelled as AI-generated	Not assessed

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	Not assessed
MAP-1	Context: intended purpose, setting and expected users are understood and recorded	Not assessed
MAP-4	Third-party components: risks from external models, data and software are mapped	Not assessed
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

Colorado SB 26-189

Applicability not yet determined. Colorado SB 26-189: obligations mapped to this record (a mapping is not compliance). Registry revision 1, Unverified requires legal review.

Control	Title	Status
CONSEQUENTIAL_DECISION	Uses in consequential decisions meet the Act's duties (notice, impact assessment, appeal)	Not assessed (date to be confirmed)

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	9	14 (6)	0

Clause	Requirement	Status	Evidence or gap
GOVERN 2.1	Organizational Ownership & Responsibility	Gap	Organizational governance requires a designated owner team.
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 (238 chars)
MAP 1.1	Deployment Context & Affected Users (optional)	Gap	Document the functional area and the customer segments the model affects.
MAP 4.1	Source & Dependency Provenance (optional)	Met	Provenance: northwind/web/branch-faq-chatbot.tsx
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	Gap	A named, resourced owner must be accountable for responding to the model's risks.
MEASURE 3.1	Risk Tracking Mechanisms	Met	3 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	5	13 (1)	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/web/branch-faq-chatbot.tsx
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	Gap	Clause 5.3 requires assigned roles and responsibilities: assign a named owner team.
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	Draft
Organisation	Northwind Finance (sample) · Lending and operations AI
Record ID	b3e90a8c-38c7-462c-9d4e-cd21f06c26e5
Content hash (SHA-256)	520a68782b0141189f1a7a7ebe588d78db3d1f23b18ec39859121c1ac908a50e
Generated	2026-11-27T15:17:16.653Z 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