Eshalu AI GOVERNANCE

Detailed AI Governance Assessment

Control-level findings, evidence and responsibilities

Prepared exclusively for

Northgate Financial Services (SAMPLE — synthetic engagement, not a client)

Reviewer-reviewed assessment

Report version 1.0	**Completed 12 August 2026**	**Record ESH-20260812-SE0001**

3 AI systems assessed · 8 priority governance actions

Confidential — prepared exclusively for the named organisation.
Eshalu AI Governance Readiness Methodology · Not legal advice or certification.

Document control

Item	Value
Prepared for	Northgate Financial Services (SAMPLE — synthetic engagement, not a client)
Report version	1.0
Assessment completed	12 August 2026
Assessment record	ESH-20260812-SE0001
Assessment basis	Eshalu AI Governance Readiness Methodology

Executive summary

Assessment status: Reviewer review complete

Reviewer review is complete. Material governance gaps remain and should be remediated before a stronger readiness position can be stated.

This assessment covers 3 AI systems operated by Northgate Financial Services (SAMPLE — synthetic engagement, not a client). It records the current assessment position, information still required and any reviewer-confirmed actions.

Overall readiness outcome: material governance gaps identified; remediation and follow-up evidence are required before a stronger readiness position can be issued.

3	1	0	8	6	5
AI systems	Specialist attention	Incomplete info	Governance gaps	Evidence files	Provider-dependent

Scope and important limitations

This is a governance-readiness assessment. It is not legal advice, certification, conformity assessment or independent assurance. The regulatory signals shown are readiness indicators, not a legal determination. References identify the recognised basis for findings and recommendations; they do not constitute legal advice or certification.

Organisation governance posture

Assessment basis: Eshalu AI Governance Readiness Methodology.

This report reflects the respondent's information and the reviewer-completed evidence review. Material governance gaps remain open; the conclusion is a governance-readiness position, not a legal or compliance determination.

Jurisdiction context recorded by the organisation: United Kingdom, European Union. The assessment provides a cross-framework governance-readiness view; any jurisdiction-specific reference remains a screening aid, not a legal determination.

Assessment methodology

Eshalu evaluates each registered AI system against the active governed control catalogue. Applicability is determined by the system's recorded risk, lifecycle, data, exposure and operating characteristics. Evidence supplied by the respondent remains subject to reviewer confirmation. Domain statuses below are qualitative summaries derived from the underlying governed control states; no unapproved percentage score is used.

How to read the status labels

Status	Meaning
Established	Assessed with accepted evidence
Partially evidenced	Partly evidenced; more required
Evidence required	Applicable; evidence not yet supplied or reviewed

Status	Meaning
Reviewer confirmation	Awaiting reviewer confirmation
Action required	Reviewer-confirmed gap to address
Not assessed	Not yet assessed / information incomplete
Not applicable	Governed as not applicable

AI system summaries

Customer Support Copilot	Maturity: Initial
Purpose & use	Answers customer billing questions from help articles and drafts adjustments for human approval
Readiness signal	EU transparency-duty candidate - specialist review recommended
Lifecycle	live
Primary responsibility	Deployer
Provider dependency	ExampleAI Ltd | LLM API V2 | GPT-4o-class
Evidence confidence	High
Top gaps	Validation; Human oversight; Data protection
Top actions	Validation (before go-live); Human oversight (before go-live); Data protection (within 30 days)
Specialist attention	Yes — reviewer/specialist attention flagged
Recommended templates	AI Testing and Validation Plan, AI Validation Report, Provider Evidence Request, Human Oversight and Override Procedure

Fraud Triage Assistant	Maturity: Developing
Purpose & use	Ranks suspected fraud alerts for analyst review
Readiness signal	not evaluated
Lifecycle	live
Primary responsibility	Shared (Provider, Deployer)
Provider dependency	ExampleAI Ltd | Content API | GPT-4o-class
Evidence confidence	High
Top gaps	Monitoring; Traceability
Top actions	Traceability (within 30 days); Monitoring (within 60–90 days)
Specialist attention	Not flagged
Recommended templates	AI Operational Monitoring Plan, AI System Inventory, AI Logging and Traceability Standard, AI System Card

Model Access Gateway	Maturity: Initial
Purpose & use	Internal gateway routing staff prompts to a third-party model
Readiness signal	not evaluated
Lifecycle	live
Primary responsibility	Shared (Deployer, GPAI provider)
Provider dependency	ExampleAI Ltd | Foundation Model | GPT-4o-class
Evidence confidence	High
Top gaps	Change control; Validation
Top actions	Validation (within 30 days); Change control (within 60–90 days)
Specialist attention	Not flagged
Recommended templates	AI Change and Release Control, Provider Evidence Request, AI Testing and Validation Plan, AI Validation Report

Domain-by-domain readiness

Compact per-system status by governance domain. See the status legend above for the meaning of each label; per-system interpretation is in the summaries.

Customer Support Copilot

Domain	Status detail	Status
Regulatory classification	Specialist review recommended	Action required
Governance and accountability	Established	Established
Data governance and privacy	Needs attention	Action required
Human oversight and impact	Needs attention	Action required
Testing and validation	Needs attention	Action required
Operations and monitoring	Evidence follow-up required	Evidence required
Evidence and assurance readiness	Not assessed	Not assessed

Fraud Triage Assistant

Domain	Status detail	Status
Regulatory classification	Specialist review not flagged	Not assessed
Governance and accountability	Evidence follow-up required	Evidence required
Data governance and privacy	Needs attention	Action required
Human oversight and impact	Not assessed	Not assessed
Testing and validation	Not assessed	Not assessed
Operations and monitoring	Needs attention	Action required
Evidence and assurance readiness	Not assessed	Not assessed

Model Access Gateway

Domain	Status detail	Status
Regulatory classification	Specialist review not flagged	Not assessed
Governance and accountability	Needs attention	Action required
Data governance and privacy	Not assessed	Not assessed
Human oversight and impact	Not assessed	Not assessed
Testing and validation	Needs attention	Action required
Operations and monitoring	Not assessed	Not assessed
Evidence and assurance readiness	Not assessed	Not assessed

Per-system readiness

Customer Support Copilot

Readiness signal: EU transparency-duty candidate - specialist review recommended. Control maturity: Initial. Evidence position: High.

Per-system governed readiness status: material gaps identified. Evidence completeness: 100%.

Risk-to-control traceability: Applicable from recorded system risk, data and operating characteristics. -> Validation (further work required); Applicable from recorded system risk, data and operating characteristics. -> Human oversight (further work required); Applicable from recorded system risk, data and operating characteristics. -> Data protection (alignment issue requiring review); Applicable from recorded system risk, data and operating characteristics. -> Accountability (established); Applicable from recorded system risk, data and operating characteristics. -> Monitoring ( ); Applicable from recorded system risk, data and operating characteristics. -> Traceability (partially established).

Responsibility context: Deployer responsibility. Provider/platform dependency: ExampleAI Ltd | LLM API V2 | GPT-4o-class.

Confirmed strengths: Accountability.

Monitoring evidence remains incomplete; follow-up evidence is required.

Areas to address:

- **Validation:** validation needs attention
- **Human oversight:** meaningful human review with real override authority should be demonstrated
- **Data protection:** data protection needs attention

Fraud Triage Assistant

Readiness signal: not evaluated. Control maturity: Developing. Evidence position: High.

Per-system governed readiness status: material gaps identified. Evidence completeness: 100%.

Risk-to-control traceability: Applicable from recorded system risk, data and operating characteristics. -> Monitoring (further work required); Applicable from recorded system risk, data and operating characteristics. -> Traceability (alignment issue requiring review); Applicable from recorded system risk, data and operating characteristics. -> Accountability (established); Applicable from recorded system risk, data and operating characteristics. -> Change control (evidence follow-up required).

Responsibility context: Shared provider and deployer responsibility. Provider/platform dependency: ExampleAI Ltd | Content API | GPT-4o-class.

Confirmed strengths: Accountability.

Change-control evidence remains incomplete; follow-up evidence is required.

Areas to address:

- **Monitoring:** ongoing monitoring with indicators and escalation should be established
- **Traceability:** logging sufficient to reconstruct decisions should be retained

Model Access Gateway

Readiness signal: not evaluated. Control maturity: Initial. Evidence position: High.

Per-system governed readiness status: material gaps identified. Evidence completeness: 100%.

Risk-to-control traceability: Applicable from recorded system risk, data and operating characteristics. -> Change control (further work required); Applicable from recorded system risk, data and operating characteristics. -> Validation (alignment issue requiring review); Applicable from recorded system risk, data and operating characteristics. -> Accountability (established).

Responsibility context: Shared provider and deployer responsibility. Provider/platform dependency: ExampleAI Ltd | Foundation Model | GPT-4o-class.

Confirmed strengths: Accountability.

Areas to address:

- **Change control:** a process to assess and approve changes should be in place
- **Validation:** validation needs attention

Prioritised findings and practical next steps

Customer Support Copilot — Validation

What: validation needs attention

Why it matters: validation gives defensible confidence the system performs as intended

Action: Define realistic tests, edge and failure cases, acceptance criteria and retained results for Customer Support Copilot before deployment and material changes.

Expected evidence: Pre-deployment and change-validation plan, realistic test cases, acceptance criteria and results.

Suggested owner: Model or quality-assurance owner. Responsibility: Deployer responsibility.

Dependency note: May require model, platform or supplier information in addition to organisation-owned evidence.

Customer Support Copilot — Human oversight

What: meaningful human review with real override authority should be demonstrated

Why it matters: effective oversight is central to safe use of higher-impact systems

Action: Define reviewer authority, intervention points, competence requirements and retained override or challenge records for Customer Support Copilot.

Expected evidence: Oversight procedure, reviewer training, intervention authority and retained override or challenge records.

Suggested owner: Business process owner. Responsibility: Deployer responsibility.

Dependency note: Primarily evidenced through the organisation's own governance records and operating practice.

Customer Support Copilot — Data protection

What: data protection needs attention

Why it matters: lawful, well-governed data handling protects people and the organisation

Action: Document the data flows, privacy assessment, applicable safeguards, access controls and retention arrangements for Customer Support Copilot.

Expected evidence: Data inventory, lawful-basis assessment, DPIA where relevant, access controls and retention record.

Suggested owner: Privacy or data-protection owner. Responsibility: Deployer responsibility.

Dependency note: May require model, platform or supplier information in addition to organisation-owned evidence.

Fraud Triage Assistant — Monitoring

What: ongoing monitoring with indicators and escalation should be established

Why it matters: problems in live AI are only caught and corrected if monitored

Action: Define monitoring indicators, review cadence, escalation criteria and retained review records for Fraud Triage Assistant.

Expected evidence: Monitoring plan, indicators, thresholds or escalation criteria, review records and issue logs.

Suggested owner: System owner. Responsibility: Shared provider and deployer responsibility.

Dependency note: Primarily evidenced through the organisation's own governance records and operating practice.

Fraud Triage Assistant — Traceability

What: logging sufficient to reconstruct decisions should be retained

Why it matters: records make decisions reconstructable and defensible

Action: Retain records sufficient to reconstruct Fraud Triage Assistant's material inputs, outputs, model/version, human review and approvals.

Expected evidence: Logs or records linking material inputs, outputs, model/version, reviewer actions and approvals.

Suggested owner: Technical owner. Responsibility: Shared provider and deployer responsibility.

Dependency note: May require model, platform or supplier information in addition to organisation-owned evidence.

Model Access Gateway — Change control

What: a process to assess and approve changes should be in place

Why it matters: uncontrolled model or supplier change can silently break governance

Action: Document and operate a change process covering model, prompt, data, workflow and supplier changes for Model Access Gateway.

Expected evidence: Change log, approval record and release procedure covering model, prompt, data, workflow and supplier changes.

Suggested owner: Product or change owner. Responsibility: Shared provider and deployer responsibility.

Dependency note: May require model, platform or supplier information in addition to organisation-owned evidence.

Model Access Gateway — Validation

What: validation needs attention

Why it matters: validation gives defensible confidence the system performs as intended

Action: Define realistic tests, edge and failure cases, acceptance criteria and retained results for Model Access Gateway before deployment and material changes.

Expected evidence: Pre-deployment and change-validation plan, realistic test cases, acceptance criteria and results.

Suggested owner: Model or quality-assurance owner. Responsibility: Shared provider and deployer responsibility.

Dependency note: May require model, platform or supplier information in addition to organisation-owned evidence.

Accountability

What: a clearly accountable owner with documented authority is needed

Why it matters: clear ownership is the basis for every other governance control

Action: Name an accountable owner for Northgate Financial Services (SAMPLE — synthetic engagement, not a client), document their authority, responsibilities and escalation route, and provide supporting evidence.

Expected evidence: Role description, delegation or governance record showing authority and escalation responsibilities.

Suggested owner: Executive sponsor or AI system owner. Responsibility: Organisation responsibility.

Dependency note: Confirm whether supplier information is also needed.

Evidence position

6 evidence files were supplied and recorded with this assessment. The reviewer assessed their relevance, currency and adequacy. A file being supplied does not by itself satisfy a control.

Action list

Advisory action list assembled from reviewer-confirmed findings. Timing bands are suggested sequencing, not regulatory deadlines. Recommended templates are aids and do not by themselves establish compliance.

Priority	Action	Responsibility & owner	Evidence & template	Timing
Action required	**Customer Support Copilot — Validation** Validation requires action for Customer Support Copilot Why: validation gives defensible confidence the system performs as intended	**Deployer · provider-dependent** Suggested owner: Model or quality-assurance owner Owner: not yet assigned	Pre-deployment and change-validation plan, realistic test cases, acceptance criteria and results. *Evidence outstanding* Template: AI Testing and Validation Plan	Before go-live Due date: to be agreed *Open - action pending*
Action required	**Customer Support Copilot — Human oversight** Human oversight requires action for Customer Support Copilot Why: effective oversight is central to safe use of higher-impact systems	**Deployer** Suggested owner: Business process owner Owner: not yet assigned	Oversight procedure, reviewer training, intervention authority and retained override or challenge records. *Evidence outstanding* Template: Human Oversight and Override Procedure	Before go-live Due date: to be agreed *Open - action pending*
Evidence required	**Customer Support Copilot — Data protection** Data protection requires action for Customer Support Copilot Why: lawful, well-governed data handling protects people and the organisation	**Deployer · provider-dependent** Suggested owner: Privacy or data-protection owner Owner: not yet assigned	Data inventory, lawful-basis assessment, DPIA where relevant, access controls and retention record. *Evidence outstanding* Template: AI Privacy and DPIA Screening Form	Within 30 days Due date: to be agreed *Open - action pending*
Evidence required	**Fraud Triage Assistant — Traceability** Traceability requires action for Fraud Triage Assistant Why: records make decisions reconstructable and defensible	**Shared · provider-dependent** Suggested owner: Technical owner Owner: not yet assigned	Logs or records linking material inputs, outputs, model/version, reviewer actions and approvals. *Evidence outstanding* Template: AI Logging and Traceability Standard	Within 30 days Due date: to be agreed *Open - action pending*
Evidence required	**Model Access Gateway — Validation** Validation requires action for Model Access	**Shared · provider-dependent** Suggested owner: Model or	Pre-deployment and change-validation plan, realistic test	Within 30 days Due date: to be agreed

Priority	Action	Responsibility & owner	Evidence & template	Timing
	Gateway Why: validation gives defensible confidence the system performs as intended	quality-assurance owner Owner: not yet assigned	cases, acceptance criteria and results. *Evidence outstanding* Template: AI Testing and Validation Plan	*Open - action pending*
Partially evidenced	**Fraud Triage Assistant — Monitoring** Monitoring requires action for Fraud Triage Assistant Why: problems in live AI are only caught and corrected if monitored	**Shared** Suggested owner: System owner Owner: not yet assigned	Monitoring plan, indicators, thresholds or escalation criteria, review records and issue logs. *Evidence outstanding* Template: AI Operational Monitoring Plan	Within 60–90 days Due date: to be agreed *Open - action pending*
Partially evidenced	**Model Access Gateway — Change control** Change control requires action for Model Access Gateway Why: uncontrolled model or supplier change can silently break governance	**Shared · provider-dependent** Suggested owner: Product or change owner Owner: not yet assigned	Change log, approval record and release procedure covering model, prompt, data, workflow and supplier changes. *Evidence outstanding* Template: AI Change and Release Control	Within 60–90 days Due date: to be agreed *Open - action pending*
Action required	**Organisation — Accountability** Accountability requires action at organisation level Why: clear ownership is the basis for every other governance control	**Organisation** Suggested owner: Executive sponsor or AI system owner Owner: not yet assigned	Role description, delegation or governance record showing authority and escalation responsibilities. *Evidence outstanding* Template: AI Governance Charter	Timing to agree determine Due date: to be agreed *Open - action pending*

Responsibility and provider-dependency summary

Where governance responsibility sits across the open actions, so provider-dependent work is visible and not mistaken for organisation-only effort.

Responsibility band	Actions	What it means
Organisation / deployer	4	Owned and evidenced by your organisation.
Provider-dependent	5	Needs model, platform or supplier information from the provider.
Shared	4	Split between your organisation and the provider.
Unresolved responsibility	0	No unresolved responsibility recorded in this sample.

Recommended templates

Priority	Recommended template	Used for
1	**AI Testing and Validation Plan**	Validation gaps
2	**AI Validation Report**	Validation gaps
3	**Provider Evidence Request**	Validation gaps, Change control
4	**Human Oversight and Override Procedure**	Oversight and override controls
5	**AI Privacy and DPIA Screening Form**	Data protection
6	**AI Full DPIA**	Data protection
7	**AI Logging and Traceability Standard**	Traceability
8	**AI System Card**	Traceability
9	**AI Operational Monitoring Plan**	Monitoring
10	**AI System Inventory**	Monitoring
11	**AI Change and Release Control**	Change control
12	**AI Governance Charter**	Accountability
13	**Roles and Responsibilities Matrix (RACI)**	Accountability
14	**AI Use Policy**	Accountability

Templates are implementation aids and do not by themselves establish compliance.

Unanswered, incomplete and excluded areas

No system was marked not assessed in the generated result.

Questions not activated by the governed progressive-assessment rules are retained in the separate Submission Audit Record together with their activation status.

Potentially relevant regulatory references

The portal records the organisation's global operating and jurisdiction context. Specific references are displayed only where the governed catalogue and recorded facts support a screening reference. They do not determine that a law applies or that the organisation complies.

Affected area	Reference	Why relevant	Basis
Validation — Customer Support Copilot	**EU AI Act Article 15**	Potentially relevant — primary readiness reference	Screening reference
Human oversight — Customer Support Copilot	**EU AI Act Article 14**	Potentially relevant — primary readiness reference	Screening reference
Data protection — Customer Support Copilot	**UK GDPR Article 35**	Potentially relevant — contributory reference	Screening reference
Validation — Model Access Gateway	**EU AI Act Article 15**	Potentially relevant — primary readiness reference	Screening reference

All references are screening references unless explicitly confirmed by a reviewer; none is a statement of legal applicability or conformity.

Assumptions and exclusions

This report relies on the information and evidence supplied by the respondent. Evidence has not been treated as accepted merely because a file was uploaded. Regulatory references are readiness-screening references and do not determine legal applicability. Detailed answers, activation decisions, declaration details and evidence metadata are retained in the separate Submission Audit Record.

Audit record reference: ESH-20260812-SE0001.

Next steps

1. Assign owners and target dates for the reviewer-confirmed actions.
2. Supply the expected evidence and complete the remediation actions.
3. Reassess affected controls after remediation to confirm closure.