SHADOW AI
DISCOVERY · SECSTUDIO

EU AI ACT READINESS ASSESSMENT

EU AI Act — Readiness & Safety Model

Prepared for **Northwind Trading Co.**

ASSESSMENT DATE	SCOPE	CLASSIFICATION
as of 12 Aug 2026	**320 endpoints · fleet-wide**	**Confidential**
PREPARED BY	REPORT ID	REGULATION
Shadow AI Discovery	**SAD-EUACT-2026-0142**	**Reg. (EU) 2024/1689**

SECTION 1

AI system inventory & classification

58
EU AI ACT READINESS / 100
Verdict: Partially ready

11
AI systems in scope

1
High-risk (Annex III)

2
GPAI models (1 self-hosted)

27
Open conformity gaps

The sensor discovered **11 distinct AI systems** across the fleet and classified each against the EU AI Act risk tiers (Reg. (EU) 2024/1689). One system — an **automated CV-screening model used in recruitment** — falls under **Annex III high-risk** and triggers the full Article 8–15 obligation set. A self-hosted general-purpose model brings **Article 53/55 GPAI provider duties** into scope. No prohibited (Article 5) practices were detected. Overall readiness is **Partially ready (58/100)**: transparency and logging gaps dominate, and the high-risk system lacks a documented risk-management system.

DISCOVERED SYSTEMS BY EU AI ACT RISK TIER

AI SYSTEM (DISCOVERED)	ROLE	RISK CLASSIFICATION	LEGAL BASIS	KEY DUTY
CV-screening / applicant ranker HR SaaS · 4 endpoints	Deployer	High-Risk	Annex III §4(a)	Art. 9–15, 26
Customer-support LLM agent Chat widget · 22 endpoints	Deployer	Limited	Art. 50(1)	Disclose AI to user
Self-hosted GPAI model Llama-class · GPU host	Provider	GPAI	Art. 53 / 55	Tech docs, copyright, eval
Public LLM API (foundation) 3rd-party · fleet-wide	Deployer	GPAI	Art. 53 (upstream)	Rely on provider docs
AI code assistant IDE plugin · 31 endpoints	Deployer	Minimal	No mandatory duty	Voluntary code of conduct
Marketing copy / image generator SaaS · 12 endpoints	Deployer	Limited	Art. 50(2)/(4)	Mark synthetic content
Meeting transcription & summary bot SaaS · 18 endpoints	Deployer	Limited	Art. 50(1)	Notify participants
Sales-forecast / demand model Internal ML · 2 hosts	Provider	Minimal	Out of Annex III	Good-practice only
Fraud / anomaly scoring Internal ML · finance	Provider	Limited	Under review	Confirm not credit-scoring

2 further minimal-risk utilities (spam filter, grammar assistant) are inventoried but carry no mandatory obligation. Classification is provisional pending Northwind confirmation of each system's purpose; a change of intended use can re-tier a system (Art. 25 reclassification).

SECTION 2 · HIGH-RISK SYSTEM

Article 9 — Risk management system

Article 9 requires a **continuous, iterative risk-management system** maintained across the lifecycle of the high-risk CV-screening system. The table below maps each sub-requirement to evidence gathered by the sensor and Northwind's documentation.

SUB-REQUIREMENT (ART. 9)	STATUS	EVIDENCE / GAP	ACTION
9(2)(a) Risk identification Known & foreseeable risks	Partial	Vendor DPIA exists; no Northwind-specific bias/discrimination risk register for the recruitment use.	Author use-case risk register incl. protected-attribute proxy risks.
9(2)(b) Risk estimation	Gap	No documented estimation of residual risk under intended and reasonably foreseeable misuse.	Run structured risk analysis; quantify likelihood/severity.
9(2)(c) Post-market data eval	Gap	No feedback loop from post-market monitoring into the risk model.	Wire monitoring signals (Sec. 4) back into risk review.
9(2)(d) Risk-control measures	Partial	Human review of shortlists exists but is undocumented and inconsistently applied.	Formalise control & oversight (Art. 14) in SOP.
9(5) Testing for compliance	Gap	No pre-deployment bias/accuracy testing against defined metrics on record.	Define test metrics & thresholds; test before each release.
9(6) Testing throughout dev	Partial	Vendor tests model; deployer-side validation on Northwind data absent.	Validate on representative local recruitment data.
9(8) Residual-risk judgement	Gap	No sign-off that each residual risk is judged acceptable.	Add accountable-owner residual-risk acceptance record.
9(9) Vulnerable-group impact	Partial	Age/disability disparate-impact not specifically assessed.	Extend bias testing to protected groups & age <18.
72 Post-market monitoring plan	Gap	No written post-market monitoring plan for the high-risk system.	Adopt Art. 72 monitoring plan template.

Article 9 obligations for high-risk systems apply from **2 Aug 2026**. The CV-screening system is **in scope now** and currently non-conformant: **5 of 9 sub-requirements are open gaps**. The risk-management system must be documented, testing evidence produced, and residual risk formally accepted before the system can carry a conformity declaration (Art. 47) and CE marking (Art. 48).

SECTION 3 · GPAI PROVIDER

Article 55 — GPAI systemic-risk obligations

Northwind self-hosts and fine-tunes a general-purpose model. Article 53 applies to all GPAI providers; **Article 55** adds obligations where a model presents **systemic risk** (≥10^{25} FLOP training compute or Commission designation). The current model is below the presumption threshold, but Northwind is treated as a provider and Art. 55 controls are assessed as forward-looking readiness.

OBLIGATION (ART. 55)	STATUS	EVIDENCE / GAP	ACTION
55(1)(a) Model evaluation & adversarial testing State-of-art eval + red-teaming	Partial	Shadow AI red-team probes captured **prompt-injection and secret-exfiltration** against the model on 6 endpoints; no standardised eval suite or documented adversarial test protocol.	Adopt a repeatable eval + red-team protocol; log results.
55(1)(b) Systemic-risk assessment & mitigation	Gap	No documented assessment of systemic risks (misuse, CBRN, cyber-offence uplift) or mitigation record.	Produce systemic-risk assessment even if below threshold.
55(1)(c) Serious-incident tracking & reporting	Gap	No process to track serious incidents or report to the AI Office / authorities without undue delay.	Stand up incident register & AI Office reporting path.
55(1)(d) Cybersecurity protection	Partial	Model host is network-segmented, but sensor found **API keys pasted into prompts** and an unregistered MCP server exposing model access.	Rotate keys; register/patch MCP; harden model weights & access.
53(1)(a) Technical documentation	Gap	No Art. 53 technical documentation package for the self-hosted model.	Compile model card, training & eval documentation.
53(1)(c) Copyright policy	Partial	Fine-tuning data provenance not fully documented; no TDM opt-out compliance policy.	Adopt copyright/TDM policy; record data sources.
53(1)(d) Training-content summary	Gap	No public summary of training content per AI Office template.	Publish training-data summary.

The **Article 55(1)(a) adversarial-testing duty** is directly supported by Shadow AI evidence: continuous endpoint red-teaming already surfaces prompt-injection, jailbreak and secret-exfiltration attempts against the self-hosted model. This telemetry is a ready-made input to a documented evaluation protocol — the control exists in practice but is **not yet formalised, versioned, or reported**. GPAI obligations have applied since **2 Aug 2025**; providers are expected to be compliant now.

SECTION 4 · HORIZONTAL CONTROLS

Cross-cutting duties

Obligations that span the high-risk and limited-risk estate — transparency, data governance, human oversight, logging and technical documentation.

CONTROL STATUS MAP

Art. 50 — Transparency to users	Gap	Art. 50(4) — Synthetic-content marking	Partial
Art. 10 — Data & data governance	Partial	Art. 14 — Human oversight	Partial
Art. 12 — Logging & record-keeping	Gap	Annex IV — Technical documentation	Gap
Art. 13 — Instructions for use	Partial	Art. 4 — AI literacy of staff	Partial

DETAIL & EVIDENCE

DUTY	STATUS	EVIDENCE / GAP	ACTION
Art. 50(1) Transparency	Gap	Customer-support agent and transcription bot do not disclose to users that they are interacting with AI.	Add AI-disclosure notice at first interaction.
Art. 50(4) Deepfake / synthetic marking	Partial	Marketing generator output not consistently machine-readable-marked as AI-generated.	Enable provenance/watermark marking on generated media.
Art. 10 Data governance	Partial	CV-screening training data lineage & representativeness not evidenced deployer-side.	Document datasets, quality & bias checks.
Art. 14 Human oversight	Partial	Human-in-the-loop exists for hiring but no defined oversight measures, stop control or competence.	Specify oversight roles, override & escalation.
Art. 12 Logging	Gap	Automatic event logging over the high-risk system's lifecycle not enabled/retained.	Enable tamper-evident logs; set retention.
Annex IV Technical docs	Gap	No Annex IV technical documentation file assembled for the high-risk system.	Compile Annex IV pack (design, data, testing, oversight).
Art. 26 Deployer obligations	Partial	Use per provider instructions partly followed; no assigned responsible owner.	Assign deployer owner; log usage & monitoring.

SECTION 5 · ROADMAP

Gap summary & conformity roadmap

GAPS BY ARTICLE

Art. 9 risk mgmt **5** Art. 55/53 GPAI **5** Cross-cutting **4** Limited-risk transparency **3**

27 open gaps total across 11 systems: 12 gap-status (critical), 11 partial, 4 monitoring/administrative. Concentrated in the high-risk CV-screening system and the self-hosted GPAI model.

PRIORITISED REMEDIATION

1 **Stand up the Article 9 risk-management system** — 0–4 weeks

Author the risk register, run bias/accuracy testing on local data, and record residual-risk acceptance for the CV-screening system. Blocks conformity.

2 **Close GPAI cybersecurity & incident gaps** — 0–4 weeks

Rotate exposed keys, register/patch the MCP server, and formalise the adversarial-test protocol and serious-incident reporting path to the AI Office.

3 **Turn on Article 12 logging & Annex IV docs** — 2–8 weeks

Enable tamper-evident lifecycle logging and assemble the Annex IV technical documentation pack for the high-risk system.

4 **Deploy Article 50 transparency notices** — 2–6 weeks

Add AI-disclosure to the support agent and transcription bot; mark synthetic marketing media as AI-generated.

5 **Formalise human oversight & deployer ownership** — 4–10 weeks

Document Art. 14 oversight measures and assign accountable Art. 26 deployer owners per system.

PATH TO CONFORMITY

For the high-risk system, the route is: complete Art. 9–15 controls → assemble Annex IV documentation → run the Art. 43 conformity assessment (internal control, Annex VI) → draw up the Art. 47 EU declaration of conformity and affix CE marking (Art. 48) → register in the EU database (Art. 71). GPAI duties (Art. 53/55) are met through the technical documentation, copyright policy, evaluation protocol and incident process rather than CE marking.

Deadlines. GPAI provider obligations (Art. 53/55) have applied since **2 Aug 2025**. High-risk obligations including Art. 9 apply from **2 Aug 2026** — now in force for the CV-screening system. Prohibited-practice bans (Art. 5) and AI-literacy (Art. 4) have applied since **2 Feb 2025**. Executing recommendations 1–3 clears all critical gaps and lifts readiness above **80/100**.

SECTION 6 · REFERENCES

References & article citations

Each EU AI Act article cited in this assessment is explained in plain English in the Shadow AI Discovery Knowledge Base. The links below deep-link to the relevant obligation so reviewers can move from a finding to the underlying legal requirement in one click. The primary source in every case is **Regulation (EU) 2024/1689** as published in the Official Journal.

EU AI ACT ARTICLES CITED IN THIS REPORT

ARTICLE	OBLIGATION (AS CITED HERE)	KNOWLEDGE BASE REFERENCE
Art. 5	Prohibited AI practices	kb/eu-ai-act.html#article-5
Art. 9	Risk-management system (high-risk)	kb/eu-ai-act.html#article-9
Art. 10	Data & data governance	kb/eu-ai-act.html#article-10
Art. 12	Logging & record-keeping	kb/eu-ai-act.html#article-12
Art. 13	Transparency & instructions for use	kb/eu-ai-act.html#article-13
Art. 14	Human oversight	kb/eu-ai-act.html#article-14
Art. 26	Deployer obligations (& Art. 25 role flip)	kb/eu-ai-act.html#article-26
Art. 50	Transparency to users & synthetic content	kb/eu-ai-act.html#article-50
Art. 53	GPAI provider obligations	kb/eu-ai-act.html#article-53
Art. 55	GPAI systemic-risk obligations	kb/eu-ai-act.html#article-55
Art. 72	Post-market monitoring	kb/eu-ai-act.html#article-72
Annex III	High-risk use cases	kb/eu-ai-act.html#annex-iii
Annex IV	Technical documentation	kb/eu-ai-act.html#annex-iv

PRIMARY & OFFICIAL SOURCES

SOURCE	LINK
Regulation (EU) 2024/1689 Full text — EUR-Lex	eur-lex.europa.eu/eli/reg/2024/1689/oj
European AI Office GPAI oversight & coordination	digital-strategy.ec.europa.eu/en/policies/ai-office
Shadow AI Discovery — EU AI Act KB Full explainer with timeline & glossary	shadow-ai-discovery.com/kb/eu-ai-act.html

Article explanations are provided for orientation and are not a substitute for the regulation text or legal advice. Where this report's classifications or gap findings turn on the precise wording of an obligation, the authoritative source is the consolidated text of Regulation (EU) 2024/1689 on EUR-Lex.

SECTION 7

Methodology & scope

HOW THIS ASSESSMENT WAS PRODUCED

Endpoint sensors continuously discover AI usage across Northwind's 320 monitored endpoints — sanctioned and shadow LLM tools, coding agents, MCP servers, self-hosted models and exposed API keys. Discovered systems are mapped to the EU AI Act (Regulation (EU) 2024/1689) using the following steps:

1 · Discovery & inventory

Passive telemetry enumerates every AI system, its role (provider / deployer / distributor) and data flows. Each artifact is deduplicated into a system-level inventory.

2 · Risk-tier classification

Each system is tested against Art. 5 (prohibited), Annex III (high-risk), Art. 50 (transparency) and the GPAI definition to assign a tier and applicable article set.

3 · Control & evidence mapping

Obligations per tier are mapped to observed evidence (e.g. red-team telemetry, key exposure, logging state) and to Northwind-supplied documentation, then scored met / partial / gap.

4 · Readiness scoring

A 0–100 composite weights open gaps by article severity and system risk tier. High-risk and systemic-risk gaps carry the greatest weight.

SCOPE & LIMITATIONS

- Covers AI systems observable from the 320 monitored endpoints as of 12 Aug 2026. Systems running entirely off-fleet or in unmonitored environments are out of scope.
- Risk-tier classifications are provisional and depend on the **intended purpose** Northwind confirms for each system; a change of use can re-tier a system.
- Evidence reflects the state observed at assessment time; controls implemented after this date are not reflected.
- Upstream third-party foundation-model provider obligations are the provider's responsibility; Northwind's duties as deployer are assessed here.

STANDARDS REFERENCED

Reg. (EU) 2024/1689 (AI Act) — Primary	**Annex III · high-risk use cases** — Applied
Annex IV · technical documentation — Applied	**GPAI Code of Practice** — Referenced

DISCLAIMER

This document is a **readiness assessment, not legal advice**. It is prepared to help Northwind Trading Co. prioritise its EU AI Act compliance programme and does not constitute a conformity assessment under Article 43, a legal opinion, or a certification of compliance. Classifications and gap findings are technical judgements based on observed evidence and information supplied by Northwind, and should be validated with qualified legal counsel and, where required, a notified body before any declaration of conformity is made. Shadow AI Discovery accepts no liability for regulatory decisions taken solely on the basis of this report. Confidential — prepared for the named recipient; do not distribute without authorization.