

EU AI Act Readiness Assessment

Kestrel HR Ltd: AI system inventory, risk classification, impact analysis and action plan

PREPARED BY ActPrism	ASSESSMENT DATE 12 October 2026	TIER Standard Assessment	VERSION 1.0
-------------------------	------------------------------------	-----------------------------	----------------

1. Executive summary

Kestrel HR Ltd ("Kestrel") is a 75-person, UK-based recruiting software company. About 40% of its customers are employers in the EU. We assessed five AI systems based on Kestrel's intake answers of 5 October 2026 and the product documentation it supplied. **Overall readiness: Moderate.** One planned feature needs immediate attention. The main compliance workload sits on a single product feature, and its deadline is 2 December 2027.

5 AI systems assessed	1 Potential prohibited practice (planned)	1 Likely high-risk system	10 Dated actions (4 × P1)
--------------------------	--	------------------------------	------------------------------

Priority finding: "Interview Insights" (planned). The roadmap describes inferring candidates' emotions and "confidence" from video interviews. This indicates a potential prohibited practice under **Art. 5(1)(f)** (emotion recognition in the workplace, which Commission guidance reads as covering recruitment). Art. 5 has applied since 2 Feb 2025. We recommend pausing the emotion-inference scope and getting a legal opinion before any further development.

- MatchScore** (candidate ranking) is **likely high-risk** under Annex III, point 4(a). The Art. 6(3) exception does not apply because the system profiles candidates. Kestrel is the *provider*. The Arts. 9–17 obligations apply from **2 Dec 2027**.
- Ask Kestrel** (support chatbot) carries **transparency obligations under Art. 50(1)**, applicable since 2 Aug 2026. The in-app widget does not yet tell users they are talking to an AI. This is a current gap.
- JD Writer** (job description generator) needs **machine-readable marking under Art. 50(2)**. Because it was on the market before 2 Aug 2026, the deadline is **2 Dec 2026**.
- GitHub Copilot** (internal) is minimal risk. Only AI-literacy measures (Art. 4) are relevant.

Next deadlines that apply to Kestrel

Date	What	Status
Already applicable	Art. 5 prohibitions; Art. 4 AI literacy; Art. 50(1) chatbot disclosure	Partial
2 Dec 2026	Art. 50(2) marking for JD Writer (legacy grace period ends)	Gap
2 Dec 2027	High-risk requirements for MatchScore (Arts. 9–17, 43, 47–49); EU authorised representative (Art. 22)	Partial

2. Scope, method and limitations

Legal baseline. Regulation (EU) 2024/1689 (the "AI Act"), as amended by Regulation (EU) 2026/1744 (the "AI Omnibus", in force since 27 July 2026). The AI Act applies in phases. The table below shows the dates that matter for this assessment.

2 Feb 2025	2 Aug 2025	2 Aug 2026	2 Dec 2026	2 Dec 2027	2 Aug 2028
Prohibitions (Art. 5), AI literacy (Art. 4)	GPAI model rules (Arts. 51–55)	General application incl. Art. 50 transparency	Art. 50(2) legacy grace ends; new Art. 5 prohibitions	High-risk: Annex III systems	High-risk: Annex I (product) systems

Teal = already applicable at assessment date. Grey = future.

Method. (1) Inventory of AI systems against the Art. 3(1) definition → (2) role per system: provider or deployer (Arts. 3, 25) → (3) prohibited-practice screen (Art. 5) → (4) high-risk screen (Art. 6(1)/Annex I, Art. 6(2)/Annex III, Art. 6(3) exception) → (5) transparency screen (Art. 50) → (6) GPAI value-chain check → (7) obligations mapping and gap analysis → (8) impact and fundamental-rights analysis.

Sources. Intake questionnaire (5 Oct 2026); product documentation v3.2; privacy policy; GDPR DPIA for MatchScore (March 2025); discovery call notes (7 Oct 2026).

Limitations. We relied on Kestrelo's answers and documents. We did not inspect code, models or data. Harmonised standards for high-risk systems are not yet finalised. Commission guidelines may change how individual provisions are interpreted.

3. AI system inventory and roles

System	Description	Model / vendor	Users	Role & rationale
S1 MatchScore	Ranks and scores applicants against a job's requirements using CV parsing and a scoring model. Recruiters see a 0–100 score and a ranked shortlist.	Proprietary ranking model + third-party LLM API for CV parsing	Customers' recruiters	Provider. Built in-house and sold under Kestrel o 's name (Art. 3(3)). Customers are deployers.
S2 Ask Kestrel o	Support chatbot on the website and in the app. Answers product questions and escalates to humans.	Third-party LLM API with retrieval	Customers' staff, prospects	Provider (own name, Art. 50(1)) and deployer on its own website.
S3 JD Writer	Generates draft job descriptions from a few inputs. Customers edit and publish them on job boards.	Third-party LLM API	Customers' recruiters	Provider of a system that generates synthetic text (Art. 50(2)).
S4 GitHub Copilot	Coding assistant used by 22 engineers.	GitHub / Microsoft	Internal engineering	Deployer. Third-party tool, used internally.
S5 Interview Insights planned	Roadmap Q1 2027: analyses recorded video interviews for "communication, confidence and emotional engagement" and produces summaries.	Not yet selected	Customers' recruiters	Provider if launched under Kestrel o 's name.

GPAI note. Kestrelo uses third-party general-purpose models through APIs and does not train or make available its own GPAI model. On the information provided, Kestrelo is **not** a GPAI model provider under Arts. 51–55. It is the provider of AI systems built on those models, and should keep the documentation its model vendors supply (see Art. 53(1)(b)).

4. Risk classification

S5 Interview Insights (planned)

Potential prohibited practice · Art. 5(1)(f)

Art. 5(1)(f) prohibits AI systems that infer the emotions of natural persons in the workplace or in education institutions, except for medical or safety reasons. The Commission's guidelines on prohibited practices read "workplace" as including candidates during recruitment. Inferring "confidence" and "emotional engagement" from facial expressions or voice falls squarely within this concern. Summarising *what* a candidate said (a transcript summary) is different and would need to be assessed on its own under Annex III, point 4(a). **Confidence: High** that the emotion component is at risk. **Recommendation:** remove emotion and affect inference from scope and get a legal opinion before development continues.

S1 MatchScore

High-risk (likely) · Annex III, point 4(a)

Art. 5: no indicators. **Annex I:** not applicable. **Annex III:** point 4(a) covers AI systems intended to be used "for the recruitment or selection of natural persons, in particular ... to analyse and filter job applications, and to evaluate candidates." MatchScore does exactly this. **Art. 6(3) exception:** not available. MatchScore evaluates candidates' suitability, which is *profiling* of natural persons, and the last subparagraph of Art. 6(3) makes profiling systems always high-risk. Its ranking also directly shapes which candidates human recruiters review, so it is not a narrow procedural or preparatory task. **Confidence: High.**

S2 Ask Kestrel

Transparency · Art. 50(1)

The chatbot interacts directly with people, so Kestrel, as provider, must design it so that users are informed they are interacting with an AI, unless that is obvious from the context. The website widget is labelled "AI assistant", but the in-app widget appears as "Kestrel Support" with a human avatar. Not high-risk: the chatbot only answers product questions and makes no decisions about people. **Applicable since 2 Aug 2026. Confidence: High.**

S3 JD Writer

Transparency · Art. 50(2)

Providers of AI systems that generate synthetic text must ensure outputs are marked in a machine-readable format and detectable as AI-generated, as far as technically feasible. The exception for systems that only perform an assistive function for standard editing is unlikely to apply, because JD Writer produces full drafts. JD Writer was on the market before 2 Aug 2026, so the deadline is 2 Dec 2026. Technical approaches (metadata, watermarking) should follow the Commission's code of practice on marking once it is final. **Confidence: Medium.** The scope for short, human-edited text is still being clarified.

S4 GitHub Copilot (internal)

Minimal risk

Internal productivity use with no decisions about people. No system-specific obligations. Kestrel should take measures to support AI literacy among staff using it (Art. 4, as amended), and its use falls under the internal AI policy. **Confidence: High.**

Classification summary

System	Classification	Role	Applies from
S5 Interview Insights	Potential prohibited	Provider (planned)	Art. 5 in force since 2 Feb 2025
S1 MatchScore	High-risk (likely)	Provider	2 Dec 2027
S2 Ask Kestrel	Art. 50(1)	Provider / deployer	2 Aug 2026
S3 JD Writer	Art. 50(2)	Provider	2 Dec 2026 (legacy grace)
S4 GitHub Copilot	Minimal	Deployer	Art. 4 since 2 Feb 2025

5. Applicable obligations and current status

Focused on **S1 MatchScore** (provider of a high-risk system), plus the transparency and general obligations already in force. Status reflects intake answers.

Obligation	Article	Applies from	Status	Observation
Risk management system	Art. 9	2 Dec 2027	Partial	Risks discussed in the DPIA. No lifecycle AI risk register.
Data and data governance	Art. 10	2 Dec 2027	Partial	Training data sources listed. No documented bias examination.
Technical documentation	Art. 11, Annex IV	2 Dec 2027	Gap	Engineering wiki only. No Annex IV structure.
Record-keeping (logs)	Art. 12	2 Dec 2027	Largely met	Scores, inputs and recruiter actions logged for 12 months.
Transparency and instructions for use	Art. 13	2 Dec 2027	Partial	Help-centre article exists. No accuracy metrics or limitations stated.
Human oversight	Art. 14	2 Dec 2027	Partial	Recruiters can override. The UI sorts by score by default, which encourages automation bias.
Accuracy, robustness, cybersecurity	Art. 15	2 Dec 2027	Partial	ISO 27001 certified. No subgroup accuracy testing; no prompt-injection testing of CV parsing.
QMS, conformity assessment, declaration, CE marking, registration	Arts. 16–17, 43, 47–49	2 Dec 2027	Gap	Not started. Internal-control route (Annex VI) expected for Annex III point 4.
EU authorised representative	Art. 22	2 Dec 2027	Gap	Kestrelro is UK-established. It must appoint an EU representative before placing MatchScore on the EU market as a high-risk system.
Post-market monitoring, serious incidents	Arts. 72–73	2 Dec 2027	Partial	Model drift monitored. No incident-reporting procedure.
Chatbot disclosure	Art. 50(1)	2 Aug 2026	Gap	In-app widget not disclosed as AI.
Machine-readable marking	Art. 50(2)	2 Dec 2026	Gap	JD Writer output not marked.
AI literacy measures	Art. 4	2 Feb 2025	Partial	Informal onboarding only. No record.

6. Impact and fundamental-rights analysis

Is an Art. 27 FRIA legally required of Kestrelro? No. Art. 27 applies to *deployers* that are public-law bodies or private entities providing public services, and to deployers of credit-scoring or life/health-insurance pricing systems. Kestrelro is a private provider. However, **three of Kestrelro's EU customers are municipalities**. As deployers of a high-risk system they will need a FRIA from 2 Dec 2027, and they will depend on information Kestrelro supplies under Art. 13. The analysis below is voluntary. It supports Art. 9 risk management and customer FRIAs.

System	Affected groups	Rights engaged	Likelihood × severity	Existing safeguards	Residual risk
S1 MatchScore	Job applicants (~40k/month across EU customers)	Non-discrimination (Charter Art. 21); data protection (Art. 8); access to work (Art. 15)	Medium × High	Human review of shortlists; DPIA; protected attributes excluded as inputs	Medium : proxy bias untested
S5 Interview Insights	Job applicants	Privacy; non-discrimination (disability, neurodivergence, culture); dignity	High × High	None (planned)	High
S2 Ask Kestrelro	Customer staff, prospects	Right to information	Low × Low	Human escalation	Low once disclosed
S3 JD Writer	Job seekers reading ads	Non-discrimination (biased wording)	Low × Medium	Customer edits before publishing	Low

7. Gap analysis and action plan

P1 = already applicable or due within 6 months · **P2** = required for the 2 Dec 2027 high-risk deadline · **P3** = good practice / commercial.

ID	Action	Article	Owner	Priority	Target
A1	Remove emotion/affect inference from Interview Insights scope. Get a legal opinion on the remaining design before building.	Art. 5(1)(f)	CEO + external counsel	P1	31 Oct 2026
A2	Label the in-app chatbot as AI ("AI assistant" name, no human avatar, first-message disclosure).	Art. 50(1)	Head of Product	P1	15 Nov 2026
A3	Add machine-readable AI-generated marking (e.g. metadata in export/API payloads) to JD Writer outputs. Track the Commission code of practice on marking.	Art. 50(2)	CTO	P1	1 Dec 2026
A4	Adopt an internal AI use policy and short AI-literacy training. Keep attendance records.	Art. 4	People Ops	P1	31 Jan 2027
A5	Set up an AI risk management process and risk register for MatchScore, linked to the existing DPIA.	Art. 9	CTO + DPO	P2	31 Mar 2027
A6	Run bias and subgroup-accuracy testing (incl. proxies). Document data provenance and representativeness. Add prompt-injection tests for CV parsing.	Arts. 10, 15	ML Lead	P2	30 Jun 2027
A7	Produce Annex IV technical documentation. Formalise post-market monitoring and serious-incident reporting.	Arts. 11, 72–73	ML Lead + Eng Manager	P2	30 Jun 2027
A8	Publish a deployer pack: instructions for use (accuracy, limitations, oversight), UI change to randomise or hide scores by default, template notices for informing workers and candidates, FRIA input sheet for public-sector customers.	Arts. 13, 14, 26, 27	Head of Product	P2	30 Sep 2027
A9	Appoint an EU authorised representative. Set up QMS procedures, complete internal-control conformity assessment, EU declaration of conformity, CE marking and EU database registration.	Arts. 16–17, 22, 43, 47–49	CEO + Compliance lead	P2	15 Nov 2027
A10	Publish a customer-facing AI Act position statement and standard procurement answers.	—	Sales / RevOps	P3	30 Nov 2026

Effort estimate. A1–A4 are low effort (roughly 2–4 person-weeks in total). A5–A9 are the core programme for MatchScore: roughly 4–6 person-months across engineering, product and compliance over 12 months. SME/small mid-cap relief measures under the amended Act may simplify the documentation format but not the substance.

Penalties, for context. Under Art. 99, maximum fines are €35M / 7% of worldwide turnover for prohibited practices and €15M / 3% for most other obligations. For SMEs, whichever amount is lower applies. Enforcement sits with national market-surveillance authorities.

8. Open questions and items for legal counsel

1. **Interview Insights.** Would a transcript-only summary (no audio or visual affect analysis) avoid Art. 5(1)(f)? How would it be classified under Annex III, point 4(a)? *Legal opinion recommended.*
2. **Customer configuration.** Some customers use MatchScore to rank *existing employees* for internal mobility. Confirm whether this falls under Annex III, point 4(b) (promotion, task allocation, monitoring) and adjust the intended-purpose statement.
3. **Substantial modification.** Kestrel re-trains the ranking model quarterly. Define which changes count as a "substantial modification" requiring re-assessment (Art. 43(4)).
4. **Registration.** Confirm EU database registration requirements against the consolidated text as amended by Reg. (EU) 2026/1744.
5. **UK position.** This report covers the EU AI Act only. UK customers and UK GDPR / Equality Act considerations are out of scope.
6. **Model vendor documentation.** Get and keep the technical information that the LLM API provider makes available to downstream providers.

Customer-facing position statement (draft)

"Kestrel has completed a structured assessment of its AI features against the EU AI Act (Regulation (EU) 2024/1689, as amended). Our candidate-ranking feature, MatchScore, is expected to be classified as high-risk under Annex III, and we are implementing the applicable requirements, including risk management, bias testing, technical documentation and human-oversight controls, ahead of the 2 December 2027 application date. Our AI assistant and generative features follow the Article 50 transparency rules. We will provide deployer documentation, including information to support customers' own impact assessments."

9. How this assessment was produced

ActPrism uses a fixed methodology mapped to the AI Act. Client answers are collected through a structured intake. A first draft is generated with AI assistance and then reviewed and edited by Eric Phan before delivery. Every classification gives its reasoning and confidence level, so your team or counsel can test it.

10. Disclaimer

This assessment was prepared by ActPrism on the basis of information supplied by the client and reflects Regulation (EU) 2024/1689, as amended by Regulation (EU) 2026/1744, and publicly available guidance as of the assessment date. It is a regulatory-readiness assessment for informational purposes only. It is **not legal advice**, does not create a lawyer–client relationship, and should not be relied on as a substitute for advice from qualified legal counsel. Conclusions depend on the accuracy and completeness of the information provided. Implementing acts, harmonised standards and Commission guidelines may affect how the obligations apply.

This is a sample. Kestrel HR Ltd, its products, customers and figures are fictional and were created to illustrate the report format. Any resemblance to a real company is coincidental.

Want this for your product?

Fixed price, delivered in 2–5 business days. Risk Snapshot \$500 · Standard Assessment \$1,200 · Full Assessment \$2,000

Book a 15-minute call: cal.com/ericphan/15min · eric@actprism.com · actprism.com