

— AI ACT · REGULATION (EU) 2024/1689

AI Act Technical Report HR & Recruitment AI

AI system inventory, preliminary risk classification, applicable obligations, gaps and an answer matrix for RFPs and procurement.

PREPARED FOR

Veltia Talent, S.L.

SYSTEMS ASSESSED

3 AI systems

DELIVERED

July 2026

METHODOLOGY

Normative tree v0.1

REGULATORY CUT-OFF

1 July 2026 (post-Omnibus)

REVIEW

Human review · v1.0

Fictional company — sample report for demonstration purposes

1 Executive summary

Veltia Talent is a recruitment SaaS selling to mid-market and enterprise clients in Spain and the EU. It uses AI at three points in its product: an in-house candidate matching engine, an in-house automated CV screening module, and an interview copilot based on a third-party provider's API, offered under the Veltia brand.

Main conclusion. The **automated CV screening** module falls, with high probability, under **Annex III 4(a)** of the Regulation (filtering and evaluation of applications): a high-risk system, with Veltia as **provider**. Full high-risk obligations apply from **2 December 2027**; the 2026–2027 window is the preparation period. The **Art. 50** transparency obligations already apply (August 2026) to the interview copilot.

Position by system

SYSTEM	ORIGIN	VELTIA'S ROLE	PRELIMINARY CLASSIFICATION
VeltiaMatch matching engine	In-house model	● PROVIDER	Likely high-risk (Annex III 4(a)), Art. 6.3 exception under analysis depending on per-client configuration
Automated CV screening	In-house module	● PROVIDER	● HIGH RISK Annex III 4(a): filtering and evaluation of applications
Interview copilot	Third-party API, Veltia brand	● DEPLOYER*	Art. 50 transparency; *own branding may convert to provider (Art. 25), gap P2

Gaps identified

6 gaps were identified (1 P1, 3 P2, 2 P3). None blocks current operations; the P1 (AI notice to candidates in the referral channel, Art. 50) is low-effort and should close within 30 days. Detail in section 6, closure plan in section 7.

What this report enables

- **Answering RFPs and procurement questionnaires** with the matrix in section 2: drafted answers, traceable to Regulation articles.
- **Prioritising** gap closure by commercial impact and regulatory date, not by intuition.
- **Demonstrating diligence** to enterprise clients: inventory, reasoned preliminary classification and a dated roadmap.

How to read the dates (post-Omnibus calendar, May 2026). Art. 50 (transparency): in force since 2 August 2026. Annex III (standalone high-risk, includes employment): 2 December 2027, a date the Commission treats as an adjustable upper limit. Annex I (AI embedded in regulated products): 2 August 2028. This report states the methodology version and its cut-off date; subsequent regulatory changes are not covered.

2 RFP / procurement answer matrix

The flagship deliverable. Each row: the question as procurement phrases it → the prepared answer (ready to paste, in Veltia's voice) → the evidence that supports it. Answers reflect the real state including gaps: in procurement, a documented "in progress" is more credible than a hollow "yes".

"Do you use AI to filter, score or rank candidates?"

● ANSWER READY

Yes. We use AI at three points: candidate-to-vacancy matching (in-house model), automated CV pre-selection (in-house module) and interview assistance (copilot on a third-party API). We maintain a formal inventory of the three systems with their preliminary risk classification under Regulation (EU) 2024/1689.

EVIDENCE · Inventory §3 · Classification §4

"Have you classified your systems under Annex III of the AI Act?"

● ANSWER READY

Yes, with a traceable methodology citing Regulation (EU) 2024/1689. The CV screening module is preliminarily classified as high-risk (Annex III 4(a), filtering and evaluation of applications). The matching engine is under Art. 6.3 exception analysis depending on configuration. The interview copilot is subject to the Art. 50 transparency obligations.

EVIDENCE · Classification §4 · ANNEX III 4(a) · ART. 6.2–6.3

"Are you the provider or deployer of each system? Where is it documented?"

● ANSWER READY

Documented per system in our AI register: we are provider of the matching engine and the screening (in-house development) and deployer of the interview copilot (third party). We are formalising the split of obligations contractually with the copilot provider, since offering it under the Veltia brand may trigger provider obligations for us (Art. 25).

EVIDENCE · Inventory §3 · ART. 3(3)–(4) · ART. 25 · Gap G-2 (§6)

"How do you inform candidates that AI is involved?"

● GAP · CLOSING

Candidates are informed in the application flow (notice in the form and in the privacy policy) and we provide clients with a template clause for their own notices. We have identified a residual channel (referral applications) where the notice is not yet shown; its closure is scheduled within 30 days.

EVIDENCE · ART. 50 · Gap G-1 (§6) · Roadmap §7

"Is there human oversight over the system's decisions?"

● GAP · CLOSING

Yes, by design: no candidate is rejected automatically; the system ranks and recommends, and the recruiter confirms each rejection. We are formalising this practice into a documented human-oversight policy (measures, training and override logging), as preparation for the high-risk obligations.

EVIDENCE · Prep. ART. 14 · Gap G-4 (§6)

"What data does the model use and what is your GDPR basis?"

● ANSWER READY

The models process application data provided by the candidate (CV, answers) and vacancy signals. We do not use special categories of data or scraping of external sources. Processing relies on legitimate interest / pre-contractual relationship depending on the case, with DPIA-2025-03 carried out for the screening and updated after the May release. Hosting and processing in the EU; DPA available.

EVIDENCE · DPIA-2025-03 · DPA · Records of processing

"Do you have technical documentation for the system (Annex IV)?"

● IN PREPARATION

Full Annex IV technical documentation is required for high-risk systems from December 2027. We hold the base today (architecture, training data, evaluation metrics and screening limitations) and its formalisation to the Annex IV schema is in our 60/90 roadmap (§7).

EVIDENCE · Roadmap §7 · Prep. ANNEX IV

"Do you keep logs of the system's decisions?"

● ANSWER READY

Yes. Every matching recommendation and every screening pre-selection is logged with a timestamp, model version and the user who confirms the decision. Log retention is 14 months (internal policy PROC-07). We will formalise model-version logging as part of the high-risk preparation.

EVIDENCE · Prep. ART. 12 · Gap G-3 (§6)

"Do you have a dated AI Act compliance plan?"

● ANSWER READY

Yes: a 30/60/90-day roadmap with owners and closure criteria (section 7 of this report), aligned with the Regulation's calendar: transparency already covered except one channel in closure; high-risk preparation in progress with a December 2027 horizon. The plan is reviewed quarterly against regulatory changes.

EVIDENCE · Roadmap §7 · Methodology §8

"Are you AI Act certified?"

● PREPARED ANSWER

There is no official AI Act certification for this type of system today. What we provide is a gap assessment with a preliminary risk classification traceable to Regulation (EU) 2024/1689, a formal inventory and a dated closure plan: the documentation procurement teams are requesting at this stage of the Regulation.

EVIDENCE · Methodology §8 · Disclaimer §9

Recommended use. Paste the answers as-is into questionnaires and RFPs, attaching this report as an evidence annex. For questions outside the matrix, get in touch before improvising: a wrong answer in procurement is hard to retract.

3 AI system inventory

Every AI system or model that touches Veltia's product, in-house or integrated. The inventory is the foundational requirement: without it there is no defensible classification.

SYSTEM	FUNCTIONAL DESCRIPTION	ORIGIN	VELTIA ROLE	PERSONAL DATA
VeltiaMatch SYS-01	Matching engine that ranks candidates by fit to the vacancy combining CV embeddings and job requirements. The recruiter sees a ranking with an explanation of factors.	In-house model (fine-tuned on Llama 3.1 8B)	● PROVIDER	CV, work history, application answers
CV screening SYS-02	Automated pre-selection: evaluates each application against minimum requirements and weighted criteria; proposes reject/advance. Rejection requires human confirmation.	In-house module (rules + classifier)	● PROVIDER	CV, requirements, scores
Interview copilot SYS-03	Generates interview guides, summarises answers and suggests follow-up questions live. Based on a third-party LLM API, offered in the Veltia interface under its brand.	Third-party API (general LLM)	● DEPLOYER*	Interview transcripts, recruiter notes

*See §4.3: offering it under own branding may trigger provider obligations (Art. 25.1). Gap G-2.

Out of inventory

Internal productivity tools (code assistants, AI office software) are out of scope for this pack as they are not part of the product and do not process candidate data. Inventorying them in the next iteration (shadow AI) is recommended.

● ANALYST NOTE

The copilot is currently marketed on Veltia's site as “Veltia AI”. That wording is exactly what can trigger Art. 25 and reclassify you as provider. Changing it to “based on third-party technology” is free and lowers the risk while the provider contract is being negotiated.

4 Preliminary classification by system

Each classification runs the normative tree v0.1: scope (Art. 2) → prohibited practices (Art. 5) → role (Art. 3, 25) → high risk (Art. 6 + Annex III, with Art. 6.3 exceptions) → transparency (Art. 50). The path and the exact reference are shown.

4.1 · SYS-02 CV screening — ● HIGH RISK · PROVIDER

Path: AI system (Art. 3.1) ✓ → no scope exclusions (Art. 2) → no Art. 5 practices → in-house development and commercialisation under the Veltia brand → **provider (Art. 3.3)** → use: “filter job applications and evaluate candidates”, i.e. **Annex III 4(a)** ✓.

Art. 6.3 exceptions analysed and discarded: the system does not perform a merely procedural or preparatory task; its output (a rejection proposal) materially influences the selection decision. Moreover, profiling of natural persons excludes the exception.

Consequence: full high-risk provider obligations from 02/12/2027 (risk management Art. 9, data Art. 10, technical documentation Art. 11 + Annex IV, logging Art. 12, transparency to deployer Art. 13, human oversight Art. 14, accuracy and robustness Art. 15, EU registration Art. 49). The 2026-2027 window is for preparation: §7.

4.2 · SYS-01 VeltiaMatch — ● LIKELY HIGH RISK · 6.3 UNDER ANALYSIS

Path: same route up to Annex III 4(a) (matching = “publish targeted job ads / analyse and filter applications”). Whether the 6.3 exception applies depends on per-client configuration: if the ranking only *orders* without rejecting and the recruiter reviews the full list, it could be argued it “improves the result of a previously completed human activity”; if the client enables automatic threshold cut-off, the exception falls away.

Recommendation: treat as high-risk for preparation purposes (conservative criterion), document per-client configuration and submit the 6.3 analysis for expert human judgement. Marked **requires expert judgement** in the tree.

4.3 · SYS-03 Interview copilot — ● ART. 50 TRANSPARENCY · ART. 25 RISK

Path: third-party system used under Veltia's authority → **deployer (Art. 3.4)**. It does not decide on candidates (it assists the interviewer): it does not fall under Annex III 4(a)-(b) as long as it neither scores nor evaluates. Interaction with people and content generation → **Art. 50.1 and 50.2:** inform of the AI interaction and mark generated content.

Art. 25.1 alert: offering the copilot under the Veltia brand may convert Veltia into the provider of the system. Contractual and branding mitigation in progress (Gap G-2). Until resolved, publicly describe the system as “based on third-party technology”.

5 Applicable obligations and calendar

In force today (from August 2026) — Art. 50

OBLIGATION	SYSTEM	STATUS
Inform people that they are interacting with an AI system (Art. 50.1), applicable to the copilot when the candidate interacts with AI outputs	SYS-03	● GAP G-1 · referral channel
Mark artificially generated content in machine-readable format (Art. 50.2); deadline extended to 02/12/2026 for certain marking solutions	SYS-03	● COVERED
AI notice in the application flow and template clause for clients (transparency best practice + GDPR arts. 13-14)	SYS-01/02	● COVERED

Required on 02/12/2027 — Annex III high-risk (preparation in the 2026–2027 window)

OBLIGATION (PROVIDER)	REF.	VELTIA'S CURRENT BASE
Continuous risk-management system	ART. 9	Partial (informal process)
Governance of training/validation data	ART. 10	Good (DPIA + documented datasets)
Technical documentation per Annex IV	ART. 11 · ANNEX IV	Base exists, not formalised
Automatic event logging	ART. 12	Good (logs with model version)
Transparency and instructions to the deployer	ART. 13	Partial (product docs)
Effective human oversight	ART. 14	Good by design, no formal policy (G-4)
Accuracy, robustness and cybersecurity	ART. 15	Metrics exist, no declared thresholds
Registration in the EU database before placing on the market	ART. 49	Pending (2027 action)

As deployer (SYS-03, and Veltia's clients regarding SYS-01/02), Art. 26 will also apply (use per instructions, oversight, relevant input data) and, for employers, worker information. The FRIA (Art. 27) does not apply to Veltia directly but may apply to some of its public-sector clients, a selling point for the template clause.

6 Prioritised gap report

Prioritised by commercial impact (does it block a sale today?) and regulatory date. Effort estimated in person-days.

ID	GAP	PRIORITY	REF.	EFFORT	SUGGESTED OWNER
G-1	AI notice missing in the referral application channel	P1 · 30 days	ART. 50	Low (2-3 d)	Product
G-2	Copilot provider/deployer role not contractually fixed; own branding exposes to Art. 25	P2 · 60 days	ART. 25	Medium (5 d + legal)	Legal + CEO
G-3	Formal register of systems and model versions (today scattered across repos)	P2 · 60 days	PREP. ART. 12/49	Medium (4 d)	CTO
G-4	Human-oversight policy undocumented (practice exists, paper does not)	P2 · 90 days	PREP. ART. 14	Medium (3 d)	CPO
G-5	Technical documentation not mapped to the Annex IV schema	P3 · 2027	ANNEX IV	High (10+ d)	CTO
G-6	Candidate-notice template clause for clients: exists but without an EN version or a public-sector (FRIA) variant	P3 · 2027	ART. 26/27	Low (2 d)	Legal

Commercial read. No gap prevents selling today. G-1 and G-2 are the only ones a demanding procurement team can flag: G-1 closes in days and G-2 is neutralised in communication (“third-party technology”) while the contract is resolved. The rest is orderly preparation for the 2027 window.

● **ANALYST NOTE**

*Human oversight exists in practice (no rejection is automatic), but there is no override log. Without that log you cannot demonstrate in an inspection that oversight is **effective** rather than a rubber stamp. It is the cheapest gap to close and the one that most reassures an enterprise buyer.*

7 Roadmap 30 / 60 / 90

— DAYS 1-30

- Close G-1: AI notice in the referral channel (Art. 50).
- Adopt this inventory as the official register; name an owner (CTO).
- Unified external messaging: copilot = “based on third-party technology” until G-2 is closed.
- Add the RFP matrix (§2) to the sales playbook; train the sales team (1-hour session).

— DAYS 31-60

- G-2: negotiate a contractual annex with the LLM provider (split of obligations, DPA, branding).
- G-3: consolidate the register of systems and versions into a single repository with a changelog.
- Decide VeltiaMatch's default configuration (no automatic cut-off) and document the 6.3 analysis with expert judgement.

— DAYS 61-90

- G-4: human-oversight policy (measures, training, override logging).
- Annex IV plan: index of technical documentation and gaps, with a 2027 calendar (G-5).
- Quarterly review: regulatory changes (AI Office / AESIA guidance) and register update.

— ONGOING MAINTENANCE

The register, classification and RFP matrix go stale with every product release and every new Commission/AESIA guidance.

Living AI Register: monthly register maintenance, regulatory-change alerts and procurement-question support, at €390/month (design-partner terms).

● ANALYST NOTE

In the session, 2 of the 5 enterprise clients had automatic threshold cut-off enabled in VeltiaMatch without Product knowing. Blocking it by default (not just documenting it) is the difference between defending the 6.3 exception and losing it.

8 Methodology

The classification applies the **Trazia normative tree v0.1**, a versioned decision tree built on the consolidated text of Regulation (EU) 2024/1689 after the Digital Omnibus (May 2026), the Commission/AI Office guidance and the AESIA FAQs. Every system runs these nodes; those the tree cannot resolve deterministically are marked *requires expert judgement* and resolved through human review.

Is it an AI system?

ART. 3.1

YES

Excluded from scope? (military, R&D, personal)

ART. 2

NO

Prohibited practice?

ART. 5

NO

Role: provider or deployer? (own branding → reclassification)

ART. 3 · 25

↓

Use under **Annex III**? — employment: filter/evaluate candidates

ANNEX III 4(a)

YES

Does the 6.3 exception apply? Is there **profiling**?

ART. 6.3

↓

Transparency obligations

ART. 50

SYS-02 Screening
Profiling → exception falls away → **high risk (provider)**

SYS-01 VeltiaMatch
6.3 depends on config → **likely high risk**

SYS-03 Copilot
Not Annex III → **Art. 50 + Art. 25 alert**

9 Scope and disclaimer

Nature of the document. This report is a **technical readiness assessment**: a preliminary risk classification with inventory, gap analysis and a preparation plan. It **is not legal advice, a legal opinion or a certification**, and it does not replace the formal legal opinion that certain decisions (contracts, litigation, EU registration) may require.

Temporal limit. The conclusions reflect the regulation in force and the guidance published as of the cut-off date (01/07/2026). The AI Act framework is under active development (implementing acts, AI Office guidance, AESIA criteria); later changes require a review of the report.

Scope limit. The analysis covers the 3 systems in the inventory (§3) based on the information provided by Veltia. It does not include forensic code review, dataset auditing or bias testing, available as additional services.

Next steps. A 90-minute delivery session to walk through the report with product and legal, resolve questions and assign owners to G-1 and G-2. The inventory and matrix should be reviewed after material product changes or new Commission/AESIA guidance. Contact: contact@usetrazia.com.

10 Preparation and version control

This report is prepared by the Trazia team from a structured inventory session, with human review before delivery. No classification is generated automatically without human review; every conclusion links to a specific normative reference.

- **Tree version:** v0.1 · **Regulatory cut-off date:** 1 July 2026.
- **Sources:** 90-minute inventory session (26/06/2026) with Veltia's CTO and DPO, product documentation provided, and verification of the public application flows.

PREPARED BY

David Cedillo, Trazia

Signature / date

LEGAL VALIDATION

Client legal counsel

Only where a conclusion requires formal legal advice

Pending where required

Version control

VERSION	DATE	CHANGE	AUTHOR
v0.9	28/06/2026	Draft after the inventory session; preliminary classification of the 3 systems.	[initials]
v1.0	04/07/2026	Methodology review, adjustment of VeltiaMatch's 6.3 analysis and incorporation of comments.	TZ

