

Risk Classification Memo — EU AI Act

Company	Aurora Hiring GmbH
AI system	Aurora CV Ranking Engine — v2.1
Date	2026-07-11
Prepared with	TrustPacket Doc Pack (self-assessment basis)

1. Conclusion

Based on the self-assessment answers, **Aurora CV Ranking Engine** appears to fall in the tier: **High risk (Annex III)**.

Your system appears to match one or more Annex III high-risk use cases. European Commission implementation material currently indicates the latest application date for Annex III high-risk rules is 2 December 2027; product-embedded Annex I systems are 2 August 2028. This is the heaviest compliance tier short of prohibition - start documentation now. Note: narrow, purely preparatory or procedural uses within these areas may be exempt under Art. 6(3) - but claiming that exemption requires a documented assessment.

2. System description

An AI system that analyzes job applications (CVs, cover letters, work samples) and produces a ranked shortlist for recruiters. The system uses natural language processing and a proprietary scoring model trained on historical hiring outcomes.

3. Intended purpose

To assist recruiters in identifying the most qualified candidates for open positions by automatically ranking applications according to job requirements, reducing time-to-hire and improving hiring quality.

4. Reasoning

Employment & HR: recruitment, screening, promotion, task allocation, monitoring

Annex III / Annex I areas reported

Employment & HR: recruitment, screening, promotion, task allocation, monitoring

5. Role and market status

Declared operator role: **Provider**. EU market status: **on the EU market / output used in the EU**.

6. Applicable obligations (summary)

Risk management system (Art. 9) — 2 Dec 2027

Data & data governance (Art. 10) — 2 Dec 2027

Technical documentation (Annex IV) (Art. 11 + Annex IV) — 2 Dec 2027

Record-keeping / logging (Art. 12) — 2 Dec 2027

Transparency & instructions for use (Art. 13) — 2 Dec 2027

Human oversight (Art. 14) — 2 Dec 2027

Accuracy, robustness & cybersecurity (Art. 15) — 2 Dec 2027

Conformity assessment + CE marking + EU registration (Art. 43, 47-49) — 2 Dec 2027

AI literacy for staff (Art. 4) — In force since 2 February 2025

This document was generated from a structured self-assessment questionnaire completed by the company named above, based on Regulation (EU) 2024/1689 (the EU AI Act). It is an informational template — not legal advice, not a conformity assessment, and not a certification. Sections marked "[To be completed]" require input the questionnaire could not capture. Review with qualified counsel before relying on it.

Technical Documentation (Annex IV structure)

Company	Aurora Hiring GmbH
AI system	Aurora CV Ranking Engine — v2.1
Date	2026-07-11
Prepared with	TrustPacket Doc Pack (self-assessment basis)

Structured to the nine Annex IV headings. For SMEs, Article 11(1) permits simplified form — every heading still requires an answer.

1. General description of the AI system

An AI system that analyzes job applications (CVs, cover letters, work samples) and produces a ranked shortlist for recruiters. The system uses natural language processing and a proprietary scoring model trained on historical hiring outcomes.

To assist recruiters in identifying the most qualified candidates for open positions by automatically ranking applications according to job requirements, reducing time-to-hire and improving hiring quality.

Deployment context: SaaS platform hosted in EU (Frankfurt region). Recruiters upload job descriptions and receive candidate rankings via web dashboard. Used by 40+ SME customers across Germany, Austria, and the Netherlands.

2. Elements of the system and development process

2.1 Models

The system integrates third-party models: OpenAI text-embedding-3-large for semantic matching; proprietary ranking model fine-tuned on anonymized hiring data. The company also trains or fine-tunes its own models.

2.2 Data and data governance (Art. 10)

- Personal data (names, emails, profiles)
- Employment / HR data (CVs, performance records)
- Behavioural / usage data (clicks, device signals)

Data sources and provenance: Training: 50,000 anonymized historical job applications (CV text, hiring outcome labels) collected from consenting customers 2020–2024. Inference: customer-uploaded CVs and job descriptions.

2.3 Human oversight measures (Art. 14)

- A human reviews or can override outputs before they take effect
- Low-confidence outputs are flagged for human attention
- Staff operating the system receive training on its limits

3. Monitoring, functioning and control

Capabilities and limitations communicated to users via instructions for use.

The system is less reliable for scanned CVs with poor OCR quality, unconventional career paths, sparse employment histories, and applications submitted in languages outside the supported set. Scores are advisory only and must not be used as the sole basis for rejecting a candidate.

4. Performance metrics

Pre-release evaluation suite / benchmark testing

Bias or fairness testing across user groups

Regression testing when models or prompts change

Latest pre-release evaluation used a holdout set of 4,200 applications across engineering, sales, and operations roles. Ranking agreement with senior recruiter panels was 0.81 NDCG@10 overall; adverse-impact ratio checks were reviewed for gender and age proxies, with any group below 0.8 routed to manual investigation before release.

5. Risk management system (Art. 9)

As a high-risk system, a documented, iterative risk management process is required across the lifecycle: identify risks, estimate and evaluate, adopt mitigation, test effectiveness.

Incident handling: AI incidents are triaged by the compliance owner within one business day. Potential discriminatory outcomes, data leakage, or recurring ranking errors are escalated to legal and engineering, with customer notification and serious-incident reporting assessment documented in the incident log.

6. Lifecycle changes

Material changes to ranking prompts, model versions, scoring thresholds, or input data categories require a change ticket, regression testing, bias review, approval by the compliance owner, and an updated release note before production deployment.

7. Standards applied

Controls are mapped to SOC 2 security and confidentiality criteria, ISO 27001 information-security controls, GDPR DPIA records, and an internal AI management procedure aligned with ISO/IEC 42001 principles. Harmonised AI Act standards will be mapped when finalized.

8. EU declaration of conformity

[Placeholder: EU declaration of conformity is finalized after the applicable conformity assessment route is complete (Art. 43, 47). Keep the supporting technical documentation available before assessment.]

9. Post-market monitoring plan (Art. 72)

Record-keeping status: The system automatically records events (logs) during operation. Log retention and access controls are managed by the provider.

The team reviews ranking quality, user override rates, appeal outcomes, drift indicators, and complaint trends monthly. Escalation is triggered by statistically significant group-performance gaps, repeated recruiter overrides, data-quality incidents, or customer reports of unfair outcomes.

This document was generated from a structured self-assessment questionnaire completed by the company named above, based on Regulation (EU) 2024/1689 (the EU AI Act). It is an informational template — not legal advice, not a conformity assessment, and not a certification. Sections marked "[To be completed]" require input the questionnaire could not capture. Review with qualified counsel before relying on it.

Model / System Card

Company	Aurora Hiring GmbH
AI system	Aurora CV Ranking Engine — v2.1
Date	2026-07-11
Prepared with	TrustPacket Doc Pack (self-assessment basis)

Overview

An AI system that analyzes job applications (CVs, cover letters, work samples) and produces a ranked shortlist for recruiters. The system uses natural language processing and a proprietary scoring model trained on historical hiring outcomes.

Intended use

To assist recruiters in identifying the most qualified candidates for open positions by automatically ranking applications according to job requirements, reducing time-to-hire and improving hiring quality.

Out-of-scope uses

The system is not designed for automated rejection, salary determination, immigration eligibility, background checks, emotion recognition, or assessing disability, health status, union membership, ethnicity, religion, or other protected characteristics.

Models and architecture

Built on third-party models: OpenAI text-embedding-3-large for semantic matching; proprietary ranking model fine-tuned on anonymized hiring data.

Training data overview

- Personal data (names, emails, profiles)
- Employment / HR data (CVs, performance records)
- Behavioural / usage data (clicks, device signals)

Training: 50,000 anonymized historical job applications (CV text, hiring outcome labels) collected from consenting customers 2020–2024. Inference: customer-uploaded CVs and job descriptions.

Evaluation

- Pre-release evaluation suite / benchmark testing
- Bias or fairness testing across user groups
- Regression testing when models or prompts change

Latest pre-release evaluation used a holdout set of 4,200 applications across engineering, sales, and operations roles. Ranking agreement with senior recruiter panels was 0.81 NDCG@10 overall; adverse-impact ratio checks were reviewed for gender and age proxies, with any group below 0.8 routed to manual investigation before release.

Limitations

The system is less reliable for scanned CVs with poor OCR quality, unconventional career paths, sparse employment histories, and applications submitted in languages outside the supported set. Scores are advisory only and must not be used as the sole basis for rejecting a candidate.

This document was generated from a structured self-assessment questionnaire completed by the company named above, based on Regulation (EU) 2024/1689 (the EU AI Act). It is an informational template — not legal advice, not a conformity assessment, and not a certification. Sections marked "[To be completed]" require input the questionnaire could not capture. Review with qualified counsel before relying on it.

Internal AI Usage Policy

Company	Aurora Hiring GmbH
AI system	Aurora CV Ranking Engine — v2.1
Date	2026-07-11
Prepared with	TrustPacket Doc Pack (self-assessment basis)

supports Article 4 AI-literacy measures: staff who operate AI systems need sufficient, role-appropriate understanding of their function, limits, context of use, and affected groups. A policy alone is not proof of compliance.

1. Purpose and scope

This policy governs how staff of Aurora Hiring GmbH use AI systems in their work, including Aurora CV Ranking Engine and third-party AI tools.

2. Approved uses

Using approved AI systems for their documented intended purpose.

Using AI assistance for drafting, analysis and internal productivity, with human review of outputs before external use.

3. Prohibited uses

Entering special-category personal data or customer confidential data into unapproved AI tools.

Using AI outputs for consequential decisions about people without the documented human oversight process.

Any use matching Article 5 prohibited practices (manipulation, social scoring, emotion recognition of staff).

4. Human oversight

A human reviews or can override outputs before they take effect

Low-confidence outputs are flagged for human attention

Staff operating the system receive training on its limits

5. Training

Staff operating Aurora CV Ranking Engine receive onboarding on its capabilities, limitations and this policy. Records of training are kept.

6. Incidents

AI incidents are triaged by the compliance owner within one business day. Potential discriminatory outcomes, data leakage, or recurring ranking errors are escalated to legal and engineering, with customer notification and serious-incident reporting assessment documented in the incident log.

7. Review

This policy is reviewed at least annually and when the company adopts new AI systems or the regulatory landscape changes.

This document was generated from a structured self-assessment questionnaire completed by the company named above, based on Regulation (EU) 2024/1689 (the EU AI Act). It is an informational template — not legal advice, not a conformity assessment, and not a certification. Sections marked "[To be completed]" require input the questionnaire could not capture. Review with qualified counsel before relying on it.

Compliance Roadmap

Company	Aurora Hiring GmbH
AI system	Aurora CV Ranking Engine — v2.1
Date	2026-07-11
Prepared with	TrustPacket Doc Pack (self-assessment basis)

Deadline-ordered plan

When	What
Already in force (2 Feb 2025)	Article 5 prohibited practices ban; Article 4 AI literacy — use the AI usage policy as one support measure, then run role-appropriate staff training and keep evidence.
Already in force (2 Aug 2025)	GPAI model duties — not applicable based on answers (you do not provide a general-purpose model).
Now	Review the evidence now captured in the Annex IV technical documentation, close any remaining highlighted gaps, and operationalise risk management (Art. 9) and logging (Art. 12).
6 months before deadline	Accuracy/robustness evidence (Art. 15) including group-level performance; finalise human oversight design (Art. 14); prepare instructions for use (Art. 13).
Before 2 Dec 2027	Conformity assessment (Art. 43) using the applicable route: internal control for most Annex III points 2-8 systems, or notified-body/product-sector routes where required. Then EU declaration of conformity, CE marking, and EU database registration (Art. 47-49). Non-EU providers: appoint an EU authorised representative (Art. 22). Annex I product-embedded systems: 2 Aug 2028.
After launch	Post-market monitoring (Art. 72) and serious-incident reporting (Art. 73).

This document was generated from a structured self-assessment questionnaire completed by the company named above, based on Regulation (EU) 2024/1689 (the EU AI Act). It is an informational template — not legal advice, not a conformity assessment, and not a certification. Sections marked "[To be completed]" require input the questionnaire could not capture. Review with qualified counsel before relying on it.

FRIA Applicability & Support Draft

Company	Aurora Hiring GmbH
AI system	Aurora CV Ranking Engine — v2.1
Date	2026-07-11
Prepared with	TrustPacket Doc Pack (self-assessment basis)

Article 27 does not apply to every high-risk deployment. It applies to deployers that are bodies governed by public law, private entities providing public services, and deployers of Annex III point 5(b) or 5(c) systems such as creditworthiness/credit scoring and certain life or health insurance risk-assessment use cases. This Full Pack document is a scope check and support draft. It should be finalized as a formal FRIA only if the deployer's role and use case bring Article 27 into scope, or if the deployer chooses to conduct one voluntarily.

1. Article 27 applicability check

Declared operator role: **Provider**. High-risk areas selected: Employment & HR: recruitment, screening, promotion, task allocation, monitoring.

Public-law body deployer: [To be completed: yes/no]

Private entity providing public services: [To be completed: yes/no]

Annex III point 5(b)/(c) creditworthiness or insurance use case: No specific 5(b)/(c) signal was captured by the questionnaire.

Annex III point 2 exception: No Annex III point 2 signal was captured by the questionnaire.

[To be completed: conclusion on whether Article 27 applies to this deployment. Do not treat this draft as a mandatory FRIA unless that conclusion is yes.]

2. Purpose and scope of this assessment

If Article 27 applies or a voluntary assessment is desired, this draft evaluates the potential impact of **Aurora CV Ranking Engine** on the fundamental rights of individuals and groups affected by its use.

2.1 System overview

An AI system that analyzes job applications (CVs, cover letters, work samples) and produces a ranked shortlist for recruiters. The system uses natural language processing and a proprietary scoring model trained on historical hiring outcomes.

To assist recruiters in identifying the most qualified candidates for open positions by automatically ranking applications according to job requirements, reducing time-to-hire and improving hiring quality.

2.2 Deployment context

SaaS platform hosted in EU (Frankfurt region). Recruiters upload job descriptions and receive candidate rankings via web dashboard. Used by 40+ SME customers across Germany, Austria, and the Netherlands.

2.3 Duration, frequency and scope of use

The system is used daily by recruiter teams at SME customers in the EU. A typical customer processes 200-2,000 applications per month, and each deployment remains active for the duration of the customer's hiring contract.

3. Affected persons and groups

Affected persons include job applicants, current employees applying for internal roles, recruiters using the dashboard, and hiring managers relying on shortlist recommendations. Particular attention is given to applicants from underrepresented groups, older candidates, candidates with disabilities, and non-native-language CVs.

4. Fundamental rights potentially affected

The following fundamental rights under the EU Charter may be engaged by this system:

Human dignity (Art. 1 Charter): [assess whether the system decisions affect dignity]

Respect for private and family life (Art. 7 Charter): The system processes personal or special-category data.

Protection of personal data (Art. 8 Charter): Covered by GDPR compliance — cross-reference your DPIA if conducted.

Equality and non-discrimination (Art. 20-21 Charter): The main equality risk is that historical hiring data or CV language patterns could disadvantage protected groups. Privacy and data-protection risks arise from processing CVs, work history, contact details, and recruiter notes. Contestability risk arises if customers treat rankings as final decisions rather than decision-support.

Right to an effective remedy (Art. 47 Charter): Customers must provide applicants with a human contact channel for questions or contestation. Recruiters can override rankings, record the reason, and request correction or deletion of applicant data according to the customer GDPR process.

5. Risk assessment

5.1 Potential harms

Potential harms include unfair deprioritization of qualified candidates, overreliance by recruiters, privacy breaches from uploaded CVs, and failure to identify model drift after role requirements change.

5.2 Likelihood and severity

Bias-related harm is rated possible and severe before mitigation, so it receives priority controls. Privacy incidents are rated rare but severe due to personal-data sensitivity. Overreliance is rated possible and moderate, mitigated through UI wording, training, and mandatory human review.

5.3 Human oversight in place

A human reviews or can override outputs before they take effect

Low-confidence outputs are flagged for human attention

Staff operating the system receive training on its limits

6. Mitigation measures

Steps taken or planned to reduce identified risks:

Technical measures: Mitigations include human review before rejection, confidence flags, fairness checks before release, role-based access controls, audit logs, encrypted storage, customer instructions for use, and release gates for model or threshold changes.

Procedural measures: Human review of outputs before consequential decisions.

Transparency measures: Customer-facing instructions state that the system provides decision support and that humans remain responsible for hiring decisions. Customers are advised to disclose AI-assisted screening in applicant notices and provide a contestation channel.

Monitoring and auditing: The team reviews ranking quality, user override rates, appeal outcomes, drift indicators, and complaint trends monthly. Escalation is triggered by statistically significant group-performance gaps, repeated recruiter overrides, data-quality incidents, or customer reports of unfair outcomes.

7. Stakeholder consultation

The sample assumes consultation with recruiter users, a customer HR lead, and an external employment-law reviewer. Feedback led to clearer override logging, stronger applicant notice wording, and a requirement to review group-level performance before release.

8. Residual risk and decision

Residual risk remains around proxy discrimination and inconsistent customer use. The deployment decision is to proceed only with documented human review, monthly monitoring, customer training, and suspension of automated ranking if material adverse-impact indicators appear.

9. Review and updates

If Article 27 applies, review and update this assessment when relevant facts change. For voluntary assessments, use the same review triggers as a good-governance control:

When the system undergoes substantial modification (Art. 43(4))

If new risks or harms are identified during deployment

At least annually as part of post-market monitoring (Art. 72)

Date of this assessment: 2026-07-11 (draft — applicability and remaining gaps must be finalized by the deployer before relying on it).

This document was generated from a structured self-assessment questionnaire completed by the company named above, based on Regulation (EU) 2024/1689 (the EU AI Act). It is an informational template — not legal advice, not a conformity assessment, and not a certification. Sections marked "[To be completed]" require input the questionnaire could not capture. Review with qualified counsel before relying on it.

Sample generated from fictional input. Your pack is based on your real answers. Documents are informational templates, not legal advice or a conformity assessment.