WORK SAMPLE · PRACTITIONER CASE

# Operationalising Human Oversight under Articles 14 and 26 EU AI Act

*A Practitioner Case for AI-Assisted CV Screening*

AUTHOR
Niklas Per Weber

AS OF
November 2026 · Legal status as of 27 July 2026

CONTACT
kontakt@nikweber.de
nikweber.de

Work sample. Fictional company, real requirements. Not legal advice. Version of 24 November 2026.

## 1. Executive Summary

Rhön Logistik GmbH (around 2,000 employees) deploys a third-party AI tool to pre-screen CVs and rank applicants, making the company a deployer of a high-risk AI system under Article 6(2) and Annex III point 4(a) EU AI Act. Nominal human-in-the-loop review fails legal and ethical standards if time-pressured recruiters rubber-stamp outputs or lack context to challenge them. This case study presents an operational oversight model aligning Articles 14 and 26 EU AI Act with ISO/IEC 42001:2023 and German BSI IT-Grundschutz. Core safeguards block auto-rejection, replace raw rank numbers with tiered feature reasons and confidence flags, enforce blind review of low-ranked samples, and require documented human reasons for every rejection. A RACI matrix, a 15-point checklist, a cross-standard controls table, and eight KPIs translate statutory duties into auditable practice.

## 2. System and Risk Context

Rhön Logistik GmbH is a mid-sized German logistics company with roughly 2,000 employees across warehousing, dispatch, fleet maintenance, and administration. Processing around 8,000 applications across 250 vacancies annually, it procures a cloud-based CV screening module from an external vendor (the provider under Article 3(3) EU AI Act) integrated into its Applicant Tracking System (ATS). Operating the tool under its own authority without modifying its intended purpose or model architecture, Rhön Logistik GmbH acts as a deployer under Article 3(4) EU AI Act.

### Function, Process Placement, and Affected Persons

Positioned between application intake and first-round interview selection, the AI module parses uploaded CVs and certificates via optical character recognition and natural language processing. It extracts qualifications, licences, work history, and language skills, matching them against job requirement profiles to support HR recruiters in shortlisting candidates.

Affected persons are external applicants and internal candidates. Pre-screening directly affects access to employment and career progression, engaging fundamental rights to non-discrimination (Article 21 EU Charter of Fundamental Rights; German General Equal Treatment Act, AGG) and data protection (Article 8 Charter).

### Main Failure Modes

1. **Bias**: Training data or proxy features (parental or caregiving career gaps, foreign vocational credentials, postal codes, gendered phrasing) can systematically penalise protected groups under Section 1 AGG.
2. **Automation Bias**: Time-pressured recruiters treat algorithmic rankings as objective truth, accepting top-ranked profiles and dismissing lower-ranked ones without checking primary documents.

3. **Opaque Rankings**: Composite scores (such as “79/100, Rank 4”) hide criterion weighting, parser uncertainty, and whether mandatory qualifications were offset by generic keyword matches.
4. **Data Drift**: Shifts in regional labour markets, new green-logistics job titles, or unfamiliar CV layouts cause input data to diverge from the provider’s validation baseline.
5. **Context Blindness**: Statistical parsers misread non-linear career paths, relevant military or self-employed logistics experience, cross-sector transitions, and cover letter explanations.

## 3. Legal Requirements

AI systems intended for recruitment or selection of natural persons, specifically to analyse and filter job applications and evaluate candidates, are classified as high-risk under Article 6(2) in conjunction with Annex III, point 4(a) of the EU AI Act (Regulation (EU) 2024/1689). According to the current consolidated text, Annex III high-risk obligations apply from 2 December 2027, whereas AI literacy duties under Article 4 have applied since 2 February 2025.

### Article 14: Human Oversight Functional Standard

While Article 14 regulates system design by providers, it defines the functional capabilities deployers must maintain in operation to prevent or minimise risks to health, safety, or fundamental rights (Article 14(1) and (2)). Under Article 14(4)(a) to (e), oversight must enable assigned natural persons to:

- **Understand capacities and limitations (Article 14(4)(a))**: Understand system capacities and limitations and monitor operation to detect and address anomalies, dysfunctions, and unexpected performance.
- **Remain aware of automation bias (Article 14(4)(b))**: Remain aware of the tendency to automatically rely or over-rely on system output, particularly for systems providing recommendations to human decision-makers.
- **Correctly interpret output (Article 14(4)(c))**: Correctly interpret output using available interpretation tools and methods.
- **Disregard, override, or reverse (Article 14(4)(d))**: Decide in any particular situation not to use the system or to disregard, override, or reverse its output.
- **Intervene or stop (Article 14(4)(e))**: Intervene in operation or interrupt the system via a “stop” button or similar procedure that halts operation in a safe state.

### Articles 13 and 26: Provider Information and Deployer Obligations

- **Instructions for use (Article 13(1) to (3))**: The provider must supply instructions specifying intended purpose, subgroup accuracy metrics, known limitations, foreseeable fundamental rights risks, input specifications, interpretability tools (Article 13(3)(b)), and Article 14 oversight measures.
- **Use per instructions (Article 26(1))**: The deployer must implement technical and organisational measures ensuring operation accords with the Article 13 instructions.
- **Competence, training, and authority (Article 26(2))**: The deployer must assign oversight to natural persons with the necessary competence, training, authority, and support.
- **Input data control (Article 26(4))**: Where controlling input data (such as job requirement profiles), the deployer must ensure relevance and representativeness.
- **Monitoring and incident reporting (Article 26(5))**: The deployer must monitor operation against instructions and inform the provider and relevant authorities of operational risks or serious incidents.
- **Log retention (Article 26(6))**: The deployer must keep automatically generated system logs under its control for at least six months, subject to applicable data protection and employment retention rules.
- **Worker information (Article 26(7))**: Before workplace deployment, the deployer must inform workers’ representatives and affected workers, aligning in Germany with works council co-determination under Sections 87(1) No. 6 and 95 of the Works Constitution Act (BetrVG).

### Articles 4 and 86 AI Act and GDPR Article 22 Interplay

Under the amended **Article 4 AI Act** (Regulation (EU) 2026/1744), deployers must take measures to support the development of AI literacy among staff operating AI systems, tailored to their knowledge, context of use, and affected persons, without guaranteeing a specific individual level. **Article 86(1) AI Act** grants affected persons subject to a deployer decision based on Annex III high-risk AI output producing legal or similarly significant adverse effects the right to clear, meaningful explanations of the AI's role and the main elements of the decision.

This intersects directly with **Article 22(1) GDPR**, which prohibits decisions based solely on automated processing, including profiling, producing legal or similarly significant effects. In *SCHUFA* (Case C-634/21, ECLI:EU:C:2023:957), the Court of Justice of the European Union held that a probability score can itself constitute an automated decision under Article 22(1) GDPR if a third party draws strongly on it so that the score plays a determining role in the decision. Uncritical rubber-stamping by human reviewers is a practical risk factor that causes a pre-selection score to play such a determining role, rather than the holding of the judgment itself. Substantive Article 14 oversight prevents the score from playing a determining role by separating AI output from independent, documented human review, avoiding prohibited solely automated rejections under Article 22(1) GDPR and creating the audit trail needed for Article 86(1) AI Act explanations.

## 4. Roles and Responsibilities

### RACI Table

*(R = Responsible, A = Accountable, C = Consulted, I = Informed. The works council is consulted and co-signs the works agreement.)*

| Activity | HR Lead (System Owner) | Trained Reviewers (HR) | AI Governance / Compliance | ISMS Officer | Data Protection Officer | Works Council | Vendor (Provider) |
|---|---|---|---|---|---|---|---|
| Evaluate Art. 13 manual and pre-use validation | A | I | R | C | C | C | C |
| Conduct DPIA and fundamental rights risk review | A | I | R | C | C | C | I |
| Inform works council (Art. 26(7)) and conclude works agreement (Sec. 87(1) No. 6 BetrVG) | A/R | I | C | I | C | C | I |
| Deliver AI literacy and oversight training (Art. 4, 26(2)) | A | C | R | I | C | I | C |
| Execute CV screening and document human reasons | A | R | I | I | I | I | I |
| Audit blind samples and monitor monthly KPIs | A | C | R | I | I | I | I |
| Maintain log integrity and retention (Art. 26(6)) | C | I | C | A/R | C | I | R |
| Trigger stop procedure and incident reporting (Art. 14(4)(e), 26(5)) | A | R | R | C | C | I | C |
| Fulfil applicant explanation requests (Art. 86 AI Act, Art. 15 GDPR) | A | C | R | I | C | I | I |

### Concrete Meaning of "Competence, Training, and Authority" (Article 26(2))

1. **Competence**: Reviewers need HR recruitment qualifications, domain knowledge of logistics roles, and basic statistical literacy to interpret false-positive/false-negative rates, parser confidence flags, and proxy correlations.

2. **Training**: Before receiving ATS screening permissions, reviewers must pass a 12-hour certification course with a practical case exam, plus annual refreshers. Training covers: (a) Article 13 system capabilities and limitations; (b) automation bias and anchoring; (c) hands-on evaluation of test CVs with synthetic parsing errors and non-linear careers; and (d) AGG non-discrimination rules and rationale documentation.
3. **Authority and Support**: Role profiles and the works agreement grant reviewers unconditional authority to override or disregard any AI output without supervisor pre-approval or career penalty, and to pause AI ranking on anomalous vacancies. Daily screening quotas are capped to protect review time.

# 5. The Oversight Design

### Phase 1: Before Use

1. **Article 13 Verification**: AI Governance and the HR Lead verify the provider's instructions for use, subgroup validation metrics, and input specifications.
2. **Structured Criteria Setup**: HR and hiring managers define objective, job-related criteria per vacancy (Article 26(4)); unvalidated free-text prompts are disabled.
3. **No Auto-Rejection**: The ATS blocks automated rejections and candidate hiding; all applications remain visible in the recruiter queue.

### Phase 2: Per-Decision Process

1. **Reasons and Confidence Instead of Raw Ranks**: To counter numerical anchoring (Article 14(4)(b) and (c)), the UI hides percentile scores and ordinal ranks. Candidates appear in three qualitative bands (*Tier A: High Match*, *Tier B: Partial Match / Review Required*, *Tier C: Low Match*) with feature-level reasons (matched, missing, or ambiguous criteria) and a parser confidence flag (*High*, *Medium*, *Low*).
2. **Mandatory Source Check on Low Confidence**: When parser confidence is *Medium* or *Low* (such as scans or multi-column layouts), the UI masks the AI tier until the reviewer opens the original CV and records an initial assessment.
3. **Mandatory Independent Look at Low-Ranked Candidates**: Per vacancy, the ATS routes a random 15% sample of *Tier C (Low Match)* candidates into the reviewer's queue with the AI tier hidden. The reviewer evaluates the primary CV independently before the tier is revealed, catching false negatives.
4. **Documented Reason for Every Rejection Relying on AI Output**: Batch rejections are disabled. Before rejecting any candidate after AI pre-screening, the reviewer must inspect the primary CV and record a specific job-related reason; "low AI score" is not selectable.
5. **Avoiding Rubber-Stamping**: Submitting a decision with under 45 seconds of active CV viewing triggers a friction prompt requiring confirmation of the verified CV section. Overriding an AI tier requires only a quick tag and short note, ensuring overrides carry less friction than rapid confirmation.

### Phase 3: Periodic Review and Stop/Escalation Criteria

Monthly audits track override rates, blind sample false negatives, and demographic selection parity. A three-tier stop procedure operationalises Article 14(4)(e) and Article 26(5):

- **Level 1 (Case Bypass)**: Reviewers click "Bypass AI Output" on any misparsed CV for manual review.
- **Level 2 (Requisition Stop)**: Reviewers or the HR Lead suspend AI ranking on a vacancy if the 15% Tier C blind sample shows a false-negative rate above 15% or if a job criterion is systematically misparsed.
- **Level 3 (System Halt)**: The HR Lead, AI Governance, or ISMS Officer halts the system across all vacancies if monthly selection parity falls below 0.80 without objective justification, an unvalidated model update occurs, or Article 26(6) logging fails. The ATS reverts to chronological manual queues and the provider is notified under Article 26(5).

# 6. Operational Oversight Checklist

| ID | Group | Yes/No Check | Legal Basis |
|---|---|---|---|
| **G-01** | Governance | Is the CV tool registered in the AI inventory as an Annex III 4(a) high-risk system with an assigned owner and verified Art. 13 manual? | Art. 6(2), Annex III 4(a), Art. 13(3), Art. 26(1) AI Act |
| **G-02** | Governance | Have a Data Protection Impact Assessment (DPIA) and fundamental rights risk review been signed off before deployment? | Art. 35 GDPR, Art. 26(1) AI Act (Art. 27 FRIA applies only to public bodies, private providers of public services and selected sectors) |
| **G-03** | Governance | Have affected workers and the works council been informed and has a works agreement been concluded before go-live? | Art. 26(7) AI Act, Sec. 87(1) No. 6, 95 BetrVG |
| **P-01** | People | Have all active reviewers completed the 12-hour AI literacy and oversight training and passed the practical assessment? | Art. 4, Art. 26(2) AI Act |
| **P-02** | People | Do job descriptions and works agreements formally guarantee reviewers authority to override AI outputs without penalty? | Art. 14(4)(d), Art. 26(2) AI Act |
| **P-03** | People | Are daily screening quotas capped in recruiter schedules to ensure sufficient time for primary document review? | Art. 26(2) AI Act |
| **PR-01** | Process | Are vacancy requirement profiles verified for job-relatedness and absence of discriminatory proxies before use? | Art. 26(4) AI Act, Sec. 1, 7 AGG |
| **PR-02** | Process | Does every post-screening rejection require a CV-verified human reason, with an active workflow for Art. 86 explanation requests? | Art. 14(4)(c) to (d), Art. 86(1) AI Act, Art. 15, 22(1) GDPR |
| **PR-03** | Process | Is a random 15% sample of Tier C (Low Match) candidates routed for blind human review on every vacancy? | Art. 14(4)(a) to (b), Art. 26(5) AI Act |
| **T-01** | Technical | Are auto-rejection and automatic hiding of low-ranked applicants technically disabled in the ATS? | Art. 14(4)(d) AI Act, Art. 22(1) GDPR |
| **T-02** | Technical | Does the UI show qualitative tiers, feature reasons, and confidence flags (not raw ranks), mask low-confidence outputs, and enforce the 45s floor? | Art. 13(3)(b), Art. 14(4)(b) to (c) AI Act |
| **T-03** | Technical | Is a role-authorised stop button active at case, vacancy, and system levels that reverts cleanly to manual queues? | Art. 14(4)(e) AI Act |
| **M-01** | Monitoring | Are immutable logs (input hash, model version, tier, reviewer ID, review time, rationale) retained for at least 6 months? | Art. 26(6) AI Act |
| **M-02** | Monitoring | Are override rates, blind sample false-negative rates, and selection rate disparities audited monthly against thresholds? | Art. 14(4)(a), Art. 26(5) AI Act |
| **M-03** | Monitoring | Is an escalation workflow active to report serious incidents or operational risks to the provider and authorities? | Art. 26(5) AI Act |

# 7. Controls Mapping and IT-Grundschutz Transfer

| Oversight Control | EU AI Act | ISO/IEC 42001:2023 (Clauses and Annex A Control Areas) | BSI IT-Grundschutz (BSI-Standard 200-2) |
|---|---|---|---|
| **Oversight Mandate and Override Authority** | Art. 14(4)(d), Art. 26(1), Art. 26(2) | Clauses 5.2 (AI policy), 5.3 (Roles, responsibilities, authorities); Annex A: AI policy, Roles and responsibilities | Security Organisation and Personnel (structure analysis, governance roles, ORP.1 / ORP.2) |

| Oversight Control | EU AI Act | ISO/IEC 42001:2023 (Clauses and Annex A Control Areas) | BSI IT-Grundschutz (BSI-Standard 200-2) |
|---|---|---|---|
| **Pre-Use Risk and Impact Assessment** | Art. 14(1) to (2), Art. 26(1) | Clauses 6.1.2 (AI risk assessment), 6.1.4 & 8.4 (AI system impact assessment); Annex A: Impact assessment | Protection Requirements Determination (*Schutzbedarfsfeststellung*) for integrity and fundamental rights |
| **Reviewer Competence and AI Literacy Training** | Art. 4, Art. 14(4) (a) to (c), Art. 26(2) | Clauses 7.1 (Resources), 7.2 (Competence), 7.3 (Awareness); Annex A: Resources (human resources) | Qualification and Sensitisation within standard protection (ORP.3) |
| **Job Requirement Calibration and Input Control** | Art. 26(4) | Clause 8.1 (Operational planning and control); Annex A: Data for AI systems | Modelling (*Modellierung*) of application and data input parameters |
| **Tiered UI, 45s Friction Floor, and 15% Blind Review** | Art. 14(4)(a) to (d) | Clauses 8.1, 8.3 (AI risk treatment); Annex A: AI system life cycle, Use of AI systems | High protection requirements (*Schutzbedarf hoch*) safeguards for high-consequence processes |
| **Documented Rejection Rationale and Explanations** | Art. 14(4)(c) to (d), Art. 86(1) | Clauses 7.4 (Communication), 7.5 (Documented information); Annex A: Information for interested parties | Operational Traceability and Verifiable Documentation (basic and standard protection) |
| **Three-Tier Stop Procedure and Vendor Escalation** | Art. 14(4)(e), Art. 26(5) | Clauses 8.1, 10.2 (Nonconformity and corrective action); Annex A: Use of AI systems, Third-party relationships | Incident Handling and Safe Fallback Architecture (high protection requirements, DER.2.1) |
| **Log Retention and Monthly KPI Audits** | Art. 26(5), Art. 26(6) | Clauses 9.1 (Monitoring and evaluation), 9.2 (Internal audit), 10.1 (Continual improvement) | Logging (OPS.1.1.5 Protokollierung), IT-Grundschutz Check (*Soll-Ist-Vergleich*), and Continuous Improvement |

### Transferring the IT-Grundschutz Method to AI Oversight

BSI-Standard 200-2 transfers directly to AI oversight by treating AI-assisted recruitment as a critical socio-technical business process. First, **structure analysis (*Strukturanalyse*)** maps the ATS, external AI module, vacancy inputs, human decision points, and vendor interfaces as target objects. Second, **protection requirements determination (*Schutzbedarfsfeststellung*)** extends classical confidentiality, integrity, and availability criteria to decision validity and fundamental rights; because wrongful pre-screening exclusions impair career access and non-discrimination rights, the process carries **high protection requirements (*hoher Schutzbedarf*)**. Third, **modelling (*Modellierung*)** pairs baseline safeguards (Article 4 literacy training, Article 26(6) logging) with supplementary controls for high protection requirements: confidence masking, the 45-second review floor, documented human rejection reasons, and the 15% blind Tier C sample. Fourth, the **IT-Grundschutz check (*IT-Grundschutz-Check*)** verifies target-versus-actual compliance via monthly KPI audits.

## 8. Metrics and Evidence

| KPI | Definition and Evidence Source | Target / Threshold | Action Trigger |
|---|---|---|---|
| **1. Override Rate** | Share of screened candidates where the reviewer changes the AI tier assignment (*ATS decision logs*). | Operating corridor: **8% to 25%** per vacancy. | **< 5%** triggers an automation bias review of the recruiter; **> 30%** triggers a model drift and criteria review with the vendor. |
| **2. Decisions with Documented Reasons** | Share of post-screening rejections with a completed, CV-verified job-related rationale (*ATS rejection log*). | **100%** technical capture; **≥ 95%** substantive quality pass in monthly peer audits. | **< 100%** triggers an immediate IT fix; **< 95%** quality pass triggers targeted reviewer coaching. |
| **3. Review Time per Case (with Floor)** | Active CV viewing duration before decision submission (*UI telemetry*). | Median **≥ 120 seconds**; floor at **45 seconds** (**< 3%** hitting warning prompt). | Triggering the 45s prompt on **> 5%** of files in a week initiates a workload check and supervisor review. |

| KPI | Definition and Evidence Source | Target / Threshold | Action Trigger |
|---|---|---|---|
| **4. Selection Rate Disparity Across Groups** | Ratio of shortlist rates across observable cohorts such as gender, age bands, and qualification origin (*anonymised audit log*). | Adverse Impact Ratio **≥ 0.80** (Four-Fifths Rule). | **< 0.80** triggers immediate feature-weight inspection; persistent unexplained disparity triggers a **Level 3 System Halt**. |
| **5. Blind Re-Review Agreement / False-Negative Rate** | Share of the 15% blind sample of Tier C candidates rated suitable (Tier A/B) by reviewers while blinded (*blind sample log*). | Tier C false-negative rate **≤ 10%**. | **> 10%** triggers vacancy criteria recalibration; **> 15%** triggers a **Level 2 Requisition Stop** and full manual review of Tier C. |
| **6. Training Completion Rate** | Share of active ATS reviewers holding a valid (< 12 months) oversight certification (*LMS and IAM logs*). | **100%** of active users. | Any lapse (**< 100%**) automatically revokes ATS AI screening permissions via access control rules. |
| **7. Incidents and Escalations** | Count of logged parsing failures, anomaly flags, and Level 1 to 3 stops, plus triage time (*incident register*). | **100%** triaged within **2 business days**; **0** open systemic defects. | Any Level 2 or 3 stop requires formal vendor root-cause analysis and Art. 26(5) reporting assessment. |
| **8. Log Completeness Rate** | Share of screening events with complete, hash-verified logs kept for at least 6 months (*ISMS log server*). | **100%** completeness and integrity. | Any gap (**< 100%**) triggers an immediate **Level 3 System Halt** until logging integrity is restored. |

## 9. Limits and Open Questions

1. **Measuring Truly "Meaningful" Oversight**: Telemetry floors (45 seconds) and mandatory text fields block rapid clicking but cannot guarantee cognitive independence; fatigued recruiters may wait out the timer and rationalise the AI tier ex post, requiring qualitative peer audits.
2. **Reviewers' Real Authority Under Workload Pressure**: Formal override rights do not remove informal seasonal pressure from hiring managers demanding fast shortlists, making active leadership and works council backing essential.
3. **Vendor Transparency Versus Trade Secrets**: SaaS providers often withhold granular feature weights and training distributions as trade secrets despite Article 13 duties, limiting independent root-cause analysis by mid-sized deployers.
4. **Fairness Auditing Versus GDPR Data Minimisation**: Auditing parity across all Section 1 AGG grounds conflicts with Article 9 GDPR bans on collecting sensitive applicant data, restricting checks to observable proxies (age, stated gender, educational geography).
5. **The "Moral Crumple Zone" Trap**: Assigning oversight to recruiters risks shifting blame for systemic model flaws onto individuals; governance must separate case-level plausibility checks (recruiter duty) from statistical validity (provider and corporate accountability).

## 10. References


- **EU AI Act**: Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L, 2024/1689, 12 July 2024, as amended by Regulation (EU) 2026/1744.
- **Digital Omnibus on AI**: Regulation (EU) 2026/1744 of 8 July 2026, OJ L, 24 July 2026.
- **GDPR**: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation), OJ L 119, 4 May 2016.
- **CJEU Case Law**: Court of Justice of the European Union, Judgment of 7 December 2023, *OQ v Land Hessen (SCHUFA Holding AG)*, Case C-634/21, ECLI:EU:C:2023:957.
- **German Law**: Works Constitution Act (*Betriebsverfassungsgesetz*, BetrVG); General Equal Treatment Act (*Allgemeines Gleichbehandlungsgesetz*, AGG).
- **ISO/IEC 42001:2023**: ISO/IEC, *Information technology, Artificial intelligence, Management system*, Geneva, 2023.
- **BSI IT-Grundschutz**: Federal Office for Information Security (*Bundesamt für Sicherheit in der Informationstechnik*), *BSI-Standard 200-2: IT-Grundschutz Methodology* and *IT-Grundschutz Compendium*, Bonn.
- **Guidance**: Article 29 Working Party / European Data Protection Board (EDPB), *Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679* (WP251rev.01).