

EU AI Act - Annex IV technical documentation

Filled worked example: a CV-screening system (Annex III, point 4(a))

Example version 1.0 - October 2026 - Regulation (EU) 2024/1689, consolidated text in force since 27 July 2026

FICTIONAL EXAMPLE - Fictional worked example. The system, the provider and every figure are invented. Educational material by LandingRed (landingred.com) — not legal advice, not an official EU template.

How to read this example

This document shows what a complete Annex IV technical documentation file can look like, using a fictional system: a CV-screening tool that parses job applications and shows recruiters a ranked shortlist. Recruiters make every decision; the system has no reject class. The fictional provider treats the system as high-risk under Annex III, point 4(a), and documents it against all nine points of Annex IV. Each section first quotes the Annex IV wording, then shows what the fictional provider writes. Use it as a calibration aid next to the empty template from landingred.com/resources/annex-iv-template - your own file must reflect your own system, and several points demand far more depth for a real system than an example can show. For high-risk systems under Article 6(2) and Annex III these requirements apply from 2 December 2027 (Article 113, point (c)(i), as amended by Regulation (EU) 2026/1744).

Annex IV, point 1 - General description of the AI system

Annex IV wording: “A general description of the AI system including: (a) its intended purpose, the name of the provider and the version of the system reflecting its relation to previous versions; (b) how the AI system interacts with, or can be used to interact with, hardware or software, including with other AI systems, that are not part of the AI system itself, where applicable; (c) the versions of relevant software or firmware, and any requirements related to version updates; (d) the description of all the forms in which the AI system is placed on the market or put into service, such as software packages embedded into hardware, downloads, or APIs; (e) the description of the hardware on which the AI system is intended to run; (f) where the AI system is a component of products, photographs or illustrations showing external features, the marking and internal layout of those products; (g) a basic description of the user-interface provided to the deployer; (h) instructions for use for the deployer, and a basic description of the user-interface provided to the deployer, where applicable;”

What the fictional provider writes:

(a) Intended purpose, provider, version

The system analyses and filters incoming job applications for office roles in the EU and presents recruiters a ranked shortlist with a score from 0 to 100 and one of three bands: strong match, possible match, review manually. It does not reject anyone. Provider: Aurora Selezione S.r.l. (an invented company, Milan). Version 2.1; version 1 (2024) was rule-based keyword matching, version 2 (2025) introduced the ranking model, version 2.1 adds the per-score explanation panel.

(b) Interaction with other hardware or software

Deployers connect the system to their applicant-tracking system through a REST API; a certified connector exists for two common ATS products. Single sign-on via the deployer's identity provider. The system contains no other AI system and sends no data to third-party AI services.

(c) Software versions and update requirements

Current release 2.1.3 (parser 1.8, ranking model 2.1, dashboard 3.0). Quarterly maintenance releases; the deployer-facing API is versioned and changes are announced 90 days ahead. Connectors require ATS versions released after January 2024.

(d) Forms of placing on the market

One form only: a subscription web service (API plus recruiter dashboard) operated by the provider. No on-premise packages, no downloads, nothing embedded in hardware.

(e) Hardware the system runs on

Provider-operated virtual machines in two EU cloud regions; inference nodes with 8 vCPU / 32 GB RAM, no GPUs at inference. The deployer needs only a current browser.

(f) Component of products

Not applicable - the system is not a component of a physical product; no photographs or layout drawings are required.

(g) User interface provided to the deployer

A recruiter dashboard per vacancy: the ranked list, each score with its three main factors, filters, a side-by-side CV view, and a one-click override that moves any application to any band with a mandatory reason field.

(h) Instructions for use

The instructions for use (document IFU-2.1, referenced here) state the intended purpose and its limits, the deployer's obligations under Article 26, the oversight configuration, input-data specifications, expected performance ranges per role family, and the procedure for reporting malfunctions and suspected incidents to the provider.

Annex IV, point 2 - Detailed description of the elements and of the development process

Annex IV wording: "A detailed description of the elements of the AI system and of the process for its development, including: (a) the methods and steps performed for the development of the AI system, including, where relevant, recourse to pre-trained systems or tools provided by third parties and how those were used, integrated or modified by the provider; (b) the design specifications of the system, namely the general logic of the AI system and of the algorithms; the key design choices including the rationale and assumptions made, including with regard to persons or groups of persons in respect of who, the system is intended to be used; the main classification choices; what the system is designed to optimise for, and the relevance of the different parameters; the description of the expected output and output quality of the system; the decisions about any possible trade-off made regarding the technical solutions adopted to comply with the requirements set out in Chapter III, Section 2; (c) the description of the system architecture explaining how software components build on or feed into each other and integrate into the overall processing; the computational resources used to develop, train, test and validate the AI system; (d) where relevant, the data requirements in terms of datasheets describing the training methodologies and techniques and the training data sets used, including a general description of these data sets, information about their provenance, scope and main characteristics; how the data was obtained and selected; labelling procedures (e.g. for supervised learning), data cleaning methodologies (e.g. outliers detection); (e) assessment of the human oversight measures needed in accordance with Article 14, including an assessment of the technical measures needed to facilitate the interpretation of the outputs of AI systems by the deployers, in accordance with Article 13(3), point (d); (f) where applicable, a detailed description of pre-determined changes to the AI system and its performance, together with all the relevant information related to the technical solutions adopted to ensure continuous compliance of the AI system with the relevant requirements set out in Chapter III, Section 2; (g) the validation and testing procedures used, including information

about the validation and testing data used and their main characteristics; metrics used to measure accuracy, robustness and compliance with other relevant requirements set out in Chapter III, Section 2, as well as potentially discriminatory impacts; test logs and all test reports dated and signed by the responsible persons, including with regard to pre-determined changes as referred to under point (f); (h) cybersecurity measures put in place;”

What the fictional provider writes:

(a) Development methods and third-party components

Development follows a documented pipeline: data preparation, feature extraction, model training, fairness gates, release review. The CV parser builds on a pre-trained open-source language model from a third party, fine-tuned by the provider on anonymised application text; the ranking model is trained in-house. Third-party components, their licences and their modification history are listed in the component register (annexed in the real file).

(b) Design specifications

General logic: version 2.1 parses each CV into skills, qualifications, years of experience and languages; a ranking model scores each application against the criteria the recruiter sets for the vacancy. Design choices and assumptions: the system is intended for applicants to office roles in the EU; name, photo, date of birth, address and nationality are removed before scoring; CVs are assumed to be in English, Italian or German. Main classification choices: a score from 0 to 100 and one of three bands - strong match, possible match, review manually - with no reject class. Optimisation target: agreement with past recruiter shortlists; required qualifications carry the most weight; years of experience count up to ten, to limit age effects. Expected output and trade-offs: a ranked shortlist with the three factors behind each score; a simpler, explainable model was chosen over a slightly more accurate one so recruiters can check each ranking (Articles 13 and 14).

(c) System architecture and computational resources

Components in order: ingestion service, parser, feature store, ranking model service, explanation layer, dashboard. Each stage writes to an audit log keyed by application id. Training uses a provider-operated EU cloud cluster; a quarterly retraining run takes roughly 100 GPU-hours; inference runs on CPU nodes.

(d) Data requirements and datasheets

Training data sets: 120,000 applications for 1,400 office vacancies, from 2021 to 2024, in Italy, Germany and Ireland; each record pairs a parsed CV with the recruiter's shortlist decision. Provenance and selection: three employers supplied the data under written agreements; vacancies with fewer than 20 applicants and applications no recruiter reviewed were excluded. Labelling: the label is the recruiter's original shortlist decision; two trained reviewers re-checked a sample of 5,000 records and a third settled disagreements. Cleaning: duplicate applications and unreadable files removed; implausible values (such as 60 years of experience) flagged and checked by hand; career-gap fields dropped as a possible gender proxy. A datasheet per data set is kept in the data governance file, where the Article 10(2)(f) bias examination also lives.

(e) Human oversight assessment (Article 14)

The oversight model is human-in-command at vacancy level: no application is rejected by the system, every shortlist is reviewed by a recruiter, and overrides take one click plus a reason. Interpretation aids under Article 13(3)(d): the three main factors per score, confidence bands, a warning when a vacancy's applicant pool differs strongly from training data, and a monthly override report that flags vacancies with unusually low override rates (possible automation bias).

(f) Pre-determined changes

Quarterly retraining on the rolling data window is a pre-determined change, bounded by a declared performance envelope: ranking quality within two points of the validated baseline and fairness gaps within the thresholds in point 2(g). A release that leaves the envelope is not deployed; it goes through a new conformity review instead.

(g) Validation and testing

Hold-out validation splits by time and by employer, so the model is tested on vacancies and organisations it never saw. Metrics: ranking agreement with recruiter shortlists (primary), recall of eventually-hired candidates in the top band, selection-rate ratio and score-gap checks across gender and age bands on the validation set, robustness tests with malformed and adversarial CVs. Every release ships with a dated test report signed by the validation lead; reports for pre-determined changes reference the envelope in point 2(f).

(h) Cybersecurity measures (Article 15)

Encryption in transit and at rest; role-based access with least privilege; signed model artefacts and a release process that refuses unsigned models; dependency and container scanning; input sanitisation and rate limiting on the API; monitoring for data poisoning signals in the training pipeline (distribution checks against the previous window); an annual third-party penetration test; incident response runbooks shared with the post-market system in point 9.

Annex IV, point 3 - Monitoring, functioning and control

Annex IV wording: “Detailed information about the monitoring, functioning and control of the AI system, in particular with regard to: its capabilities and limitations in performance, including the degrees of accuracy for specific persons or groups of persons on which the system is intended to be used and the overall expected level of accuracy in relation to its intended purpose; the foreseeable unintended outcomes and sources of risks to health and safety, fundamental rights and discrimination in view of the intended purpose of the AI system; the human oversight measures needed in accordance with Article 14, including the technical measures put in place to facilitate the interpretation of the outputs of AI systems by the deployers; specifications on input data, as appropriate;”

What the fictional provider writes:

Capabilities and limitations

The system supports shortlisting for office roles with at least 20 applicants per vacancy. Validated ranking agreement and the top-band recall are stated per role family in the test report, together with the measured values per gender and age band; performance degrades on low-volume vacancies, on CVs outside English, Italian or German, and on role families absent from training data - the dashboard labels these cases 'review manually'.

Foreseeable unintended outcomes

Main risks: indirect discrimination through proxy variables, automation bias (recruiters following the ranking uncritically), and self-reinforcing shortlists if past decisions were themselves skewed. Mitigations: proxy screening in point 2(d), the override reporting in point 2(e), fairness gates per release, and the post-market review in point 9. No risks to physical health or safety were identified for the intended purpose.

Oversight and interpretation measures

As in point 2(e): mandatory human review, one-click override with reason, per-score factors, confidence bands, pool-shift warnings, monthly override reports.

Input data specifications

CVs as PDF or DOCX up to 10 MB, plus the structured application fields from the ATS; the vacancy's criteria form is mandatory. The IFU lists the fields the deployer must not feed into the system (photographs, free-text interview notes).

Annex IV, point 4 - Appropriateness of the performance metrics

| Annex IV wording: "A description of the appropriateness of the performance metrics for the specific AI system;"

What the fictional provider writes:

Why these metrics

The system's job is to order applications for human review, not to decide. Ranking agreement with recruiter shortlists measures exactly that job; recall of eventually-hired candidates in the top band checks that the ordering does not bury good candidates; selection-rate ratios and score gaps across gender and age bands are monitored because recruitment is where the system's fundamental-rights risk is concentrated; robustness metrics cover malformed input. Raw accuracy of a reject decision is deliberately absent: the system has no reject class.

Annex IV, point 5 - Risk management system (Article 9)

| Annex IV wording: "A detailed description of the risk management system in accordance with Article 9;"

What the fictional provider writes:

Process

A continuous, iterative process over the whole lifecycle: risk identification per release and per new deployment context, analysis and estimation in a maintained risk register, treatment through data, model and process controls, and acceptance of the documented residual risk by the provider's management before each release. The register's top entries are the risks in point 3; each carries its controls, owner, the residual-risk rating and the review date. Testing against the thresholds in point 2(g) closes the loop; results feed the next iteration and the post-market system in point 9.

Annex IV, point 6 - Lifecycle changes

| Annex IV wording: "A description of relevant changes made by the provider to the system through its lifecycle;"

What the fictional provider writes:

Change log (summary)

Version 1 (2024): rule-based keyword matching. Version 2 (2025): ranking model introduced; full re-validation; assessed as a substantial modification and taken through conformity assessment again. Version 2.1 (2026): explanation panel and override reporting added; within the pre-determined-change envelope, documented in the release file. Each change record states the date, the reason, the affected components, the re-tests run and the re-assessment conclusion.

Annex IV, point 7 - Harmonised standards or solutions adopted

Annex IV wording: “A list of the harmonised standards applied in full or in part the references of which have been published in the Official Journal of the European Union; where no such harmonised standards have been applied, a detailed description of the solutions adopted to meet the requirements set out in Chapter III, Section 2, including a list of other relevant standards and technical specifications applied;”

What the fictional provider writes:

Status in this example

At the date of this fictional example the provider applies no harmonised standard published in the Official Journal for these requirements. The file therefore describes the solutions adopted per requirement of Chapter III, Section 2 - risk management (point 5), data governance (point 2(d)), technical documentation (this file), record-keeping (the audit log in point 2(c)), transparency and instructions (point 1(h)), human oversight (point 2(e)), accuracy, robustness and cybersecurity (points 2(g), 2(h), 3) - and lists the technical specifications consulted (international AI risk-management and bias guidance documents, named with their editions in the real file). When harmonised standards are published, the provider maps this file against them and records the gaps here.

Annex IV, point 8 - EU declaration of conformity

Annex IV wording: “A copy of the EU declaration of conformity referred to in Article 47;”

What the fictional provider writes:

Where it lives in this example

The one-page declaration with the Annex V content - system name and version, provider identity, sole-responsibility statement, the conformity statement, the personal-data statement, references used, place and date of issue, signatory - is attached as the final page of this document. In a real file it is signed and kept for ten years after placing on the market (Article 16, point (g), and Article 47), and a copy accompanies the registration.

Annex IV, point 9 - Post-market monitoring (Article 72)

Annex IV wording: “A detailed description of the system in place to evaluate the AI system performance in the post-market phase in accordance with Article 72, including the post-market monitoring plan referred to in Article 72(3).”

What the fictional provider writes:

Plan summary

The post-market monitoring plan collects, per quarter and per deployer: ranking-quality indicators against the validated baseline, fairness indicators (selection-rate ratio, score gaps), the override rate per vacancy, input-pool shift warnings, and deployer feedback through a named channel in the dashboard. Thresholds mirror point 2(g); a breach opens a documented investigation, and a confirmed serious incident follows the Article 73 procedure within its deadlines. Findings feed the risk register (point 5) and the retraining envelope (point 2(f)).

EU declaration of conformity (fictional specimen - Annex V)

- 1. AI system** CV-screening system, version 2.1 (all versions 2.1.x).
- 2. Provider** Aurora Selezione S.r.l. (an invented company), Via Esempio 1, 20100 Milan, Italy.
- 3. Responsibility** This EU declaration of conformity is issued under the sole responsibility of the provider.
- 4. Conformity statement** The AI system named above is in conformity with Regulation (EU) 2024/1689.
- 5. Personal data** The AI system involves the processing of personal data and complies with Regulations (EU) 2016/679 and (EU) 2018/1725 and Directive (EU) 2016/680, to the extent applicable.
- 6. References** No harmonised standard references published at the date of issue; solutions documented in the technical documentation, point 7.
- 7. Place, date, signature** Milan, [date]. [Name, function] - FICTIONAL SPECIMEN, NOT SIGNED.