

EU AI Act Annex IV — Technical Documentation

AcmeScreen v2.4.0 — automated CV-screening module

FIKTIVT EKSEMPEL — IKKE EN RIGTIG KUNDE

Acme HR-tech ApS er en konstrueret demonstrationskunde. Alle data og bias-metrikker er illustrative, og alle tal, navne, certifikater og referencer er fiktive. Dokumentet viser format, struktur og detaljeringsgrad — ikke en faktisk overensstemmelsesvurdering. Skabelonen udleveres til virkelige Modul 1-kunder og udfyldes da med kundens egne, verificerede oplysninger. PowerQuant ApS er ikke et advokatfirma, og dokumentet er ikke juridisk rådgivning.

Customer	Acme HR-tech ApS (CVR 99999999) — fictional demonstration customer
Product / AI system	AcmeScreen — automated CV-screening module
System version	v2.4.0
Provider (Art. 3(3))	Acme HR-tech ApS — provider within the meaning of Art. 3(3); employer-customers are deployers within the meaning of Art. 3(4)
Risk classification	High-risk AI system — Annex III point 4(a) (employment: recruitment and selection, analysis and filtering of job applications)
Classification reference	Regulation (EU) 2024/1689, Annex III point 4(a), as amended by Regulation (EU) 2026/1744
Legal basis for this file	Art. 11(1) and Annex IV — standalone high-risk system (Art. 11(2) does not apply)
Obligations apply from	2027-12-02 for the Annex III and Art. 26 obligations; Art. 4 and Art. 50 obligations already apply
Conformity assessment route	Internal control (Annex VI) per Art. 43(2) — no notified body required for an Annex III point 4 system; CE marking affixed per Art. 48
EU declaration reference	EU-DOC-AcmeScreen-v2.4.0 (template, not yet issued — see §10)
Document version	2.0
Document date	2026-08-31
Supersedes	Document version 1.0 of 2026-05-07 — see §12 for the change record
Review cycle	Quarterly (next review 2026-11-30); full re-issue on substantial modification per Art. 43(4)
Prepared by	PowerQuant ApS (CVR 46274067) — Modul 1

Contents

Regulatory basis, applicability timeline and competent authorities	4
R.1 Legal instrument and the correct basis for this file	4
R.2 Applicability timeline	4
R.3 Competent authorities in Denmark	5
R.4 Administrative fines	5
§1 — General description of the AI system	6
1.1 Intended purpose	6
1.2 Provider, deployers and affected persons	6
1.3 Version and release identification	6
1.4 Hardware and software architecture	7
1.5 User interface provided to the deployer	7
1.6 Instructions for use provided to the deployer	7
§2 — Detailed description of elements and development process	9
2.1 Training data	9
2.2 Feature engineering	9
2.3 Model architecture and training	9
2.4 Third-party libraries	10
2.5 Data labelling protocol	10
2.6 Validation methodology and performance metrics	10
2.7 Appropriateness of the performance metrics	10
§3 — Monitoring, functioning and control	12
3.1 Capabilities and limitations in performance	12
3.2 Expected accuracy for specific groups of persons	12
3.3 Foreseeable unintended outcomes and sources of risk	13
3.4 Input-data specifications	13
3.5 Observability and alerting	14
3.6 Drift detection	14
3.7 Human control points in the processing flow	14
3.8 Logging	14
§4 — Risk management system	15
4.1 Identified risks	15
4.2 Risk-mitigation table	15
4.3 Continuous review	16
4.4 Documented deviation and acceptance decision — DEV-2026-014	16
4.5 Criteria for judging residual risk acceptable	17
§5 — Data and data governance	18
5.1 Training-set composition	18
5.2 Bias examination and the legal basis for it	18
5.3 Data-quality controls	19
5.4 Data-cleaning protocol and the canonical identifier list	19
5.5 Data protection impact assessment	19
5.6 Lawful basis under the GDPR	19
5.7 Special-category data	20

§6 — Human oversight	21
6.1 Oversight design and the allocation between provider and deployer	21
6.2 Measures giving effect to Article 14(4)	21
6.3 Competence and training of the overseers	22
6.4 Automation-bias controls	22
6.5 Monitoring the effectiveness of oversight	22
6.6 What the deployer must do	22
§7 — Accuracy, robustness and cybersecurity	24
7.1 Accuracy	24
7.2 Robustness and adversarial testing	24
7.3 Cybersecurity	24
7.4 Failure modes	24
§8 — Risk management summary	25
§9 — Harmonised standards and other specifications applied	26
9.1 Status of harmonised standards	26
9.2 Standards applied	26
9.3 What the standards do not cover	26
9.4 Update trigger	27
§10 — EU declaration of conformity	28
§11 — Post-market monitoring	29
11.1 Monitoring plan	29
11.2 Reporting of serious incidents	29
11.3 Escalation thresholds	30
11.4 Sanctions exposure	30
§12 — Record of changes through the lifecycle	31
12.1 Change record	31
12.2 Pre-determined changes	31
12.3 Corrections made in document version 2.0	31
Annex A — Annex IV conformity map	33
Annex B — How PowerQuant produced this report	34

Regulatory basis, applicability timeline and competent authorities

R.1 Legal instrument and the correct basis for this file

This technical documentation is drawn up under **Article 11(1)** and **Annex IV** of Regulation (EU) 2024/1689 (the AI Act), **as amended by Regulation (EU) 2026/1744** (the "Digital Omnibus on AI"), published in the Official Journal on 24 July 2026 and in force from 27 July 2026.

Article 11(1) is the operative provision: it requires the technical documentation of a high-risk AI system to be drawn up **before** the system is placed on the market or put into service, and to be kept up to date thereafter. Article 11(2) is **not** applicable here. Article 11(2) addresses providers of products covered by the Union harmonisation legislation listed in **Annex I section A**, who may draw up a single set of technical documentation covering both regimes. AcmeScreen is a **standalone** high-risk system classified under **Annex III point 4(a)**; it is not a safety component of a regulated product. Earlier versions of this document cited Art. 11(2); that citation was incorrect and has been corrected (§12).

Throughout this document, "the AI Act" means Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744. Article numbers are those of the consolidated text.

R.2 Applicability timeline

Provision	Subject	Applies from	Position of AcmeScreen
Art. 113	Entry into force of the AI Act	2024-08-01	In force
Art. 4	AI literacy of staff operating or using AI systems	2025-02-02	In force — applies today to Acme and to every deployer (§6.3)
Art. 5	Prohibited AI practices	2025-02-02	In force — screened; no prohibited practice is engaged (§4.1)
Art. 50(1), (3), (4)	Transparency: AI-interaction disclosure, emotion recognition, marking of deepfakes	2026-08-02	In force. 50(1) and 50(3) are not engaged; see §1.6
Art. 50(2)	Machine-readable marking of synthetic content by providers of generative AI	2026-08-02; systems already placed on the market before that date have until 2026-12-02	Not engaged — AcmeScreen is not a generative system
Art. 5 (new prohibitions)	AI-generated CSAM and non-consensual intimate imagery	2026-12-02	Not engaged
Annex III, Arts. 8–22 and Art. 26	Provider and deployer obligations for standalone high-risk systems	2027-12-02	Primary compliance date for this system
Art. 57	National AI regulatory sandboxes operational in each Member State	2027-08-02	Participation under consideration
Annex I	High-risk AI embedded in regulated products	2028-08-02	Not applicable

The deferral of the Annex III applicability date from 2 August 2026 to **2 December 2027** follows from Regulation (EU) 2026/1744. Two points should be read together with it, because they are frequently misunderstood:

- The amendment defers the date from which the Annex III obligations become **enforceable**. It does **not** reduce their substance, and it does not defer Art. 4 or Art. 50, which apply now.
- The technical documentation required by Art. 11(1) must exist **before** the system is placed on the market. For a system already on the market on 2 December 2027, the file must be complete on that date. Building it in the intervening period is the intended sequence, not an optional early start.

R.3 Competent authorities in Denmark

Digitaliseringsstyrelsen (the Danish Agency for Digital Government, under the Ministry of Finance) is the **central coordinating authority, the notifying authority and the national contact point** for the AI Act in Denmark, pursuant to **lov nr. 467 af 14. maj 2025**, in force **2 August 2025**. This was the first national supplementary AI act adopted in the Union. Digitaliseringsstyrelsen is a designated and operating authority; it is not a prospective or future authority, and earlier versions of this document that described it as such have been corrected (§12).

Market surveillance is split by subject matter:

Authority	Scope
Digitaliseringsstyrelsen	Central coordination, notifying authority, national contact point; market surveillance for Art. 5(1)(a)–(c), (e) and (f), including remote biometric identification
Datatilsynet	Market surveillance for Art. 5(1)(d) and (g), including biometric categorisation; and, separately, supervisory authority under the GDPR
Domstolsstyrelsen	The courts' own use of AI

The national market-surveillance authority for **Annex III high-risk systems** and for Art. 50 transparency is still being finalised through bill 2025/1 LSF 111 and a forthcoming executive order. Until it is designated, Digitaliseringsstyrelsen is the addressee for AI Act correspondence as national contact point. This section is reviewed each quarter and updated when the executive order is issued.

R.4 Administrative fines

Provi-sion	Infringement	Maximum administrative fine
Art. 99(3)	Non-compliance with the prohibited practices in Art. 5	EUR 35 000 000 or 7 % of total worldwide annual turnover for the preceding financial year, whichever is higher
Art. 99(4)	Non-compliance with obligations other than Art. 5 — including the provider obligations in Arts. 16–22, the requirements of Arts. 9–15, the deployer obligations in Art. 26 and the transparency obligations in Art. 50	EUR 15 000 000 or 3 % of total worldwide annual turnover, whichever is higher
Art. 99(5)	Supplying incorrect, incomplete or misleading information to notified bodies or national competent authorities in reply to a request	EUR 7 500 000 or 1 % of total worldwide annual turnover, whichever is higher

For SMEs, including start-ups, each of the above ceilings applies as the **lower** of the two figures rather than the higher (Art. 99(6)). Fines are imposed by the national market-surveillance authority; the Art. 99 ceilings are maxima, and Art. 99(7) sets out the factors to be weighed in each case.

§1 — General description of the AI system

Annex IV point 1. Cross-references: Art. 11(1); Art. 13 for the instructions for use.

1.1 Intended purpose

The **intended purpose** of AcmeScreen, within the meaning of Art. 3(12), is to perform an *automated initial screening of job applications*: ranking inbound CVs against a job description supplied by the employer-customer and surfacing the top N candidates (default N = 20) for review by a human recruiter.

The system is **not** intended, and is not fit, to:

- make or communicate a final hiring decision;
- reject a candidate without human review — rejection always requires an explicit, authenticated action by a trained recruiter (§6.1);
- assess personality, character, attitude or emotional state; the system performs no emotion recognition within the meaning of Art. 3(39), so Art. 50(3) is not engaged;
- evaluate candidates for promotion, internal mobility or termination — these uses are contractually excluded in the data processing agreement and technically blocked, because every request must carry a valid external job-posting identifier;
- verify the truth of any claim made in a CV.

Reasonably foreseeable misuse within the meaning of Art. 3(13) is addressed in §3.3 and mitigated through the instructions for use (§1.6) and the deployer obligations flowed down in the DPA (§6.6).

1.2 Provider, deployers and affected persons

Role	Party	Basis
Provider	Acme HR-tech ApS	Art. 3(3) — develops the system and places it on the market under its own name
Deployer	The employer-customer using AcmeScreen to screen its own applicants	Art. 3(4) — uses the system under its own authority
Affected person	The job applicant	Art. 26(11), Art. 86
Operator	The recruiter, employed by the deployer	Assigned human oversight under Art. 14 and Art. 26(2)

- **Geographic markets:** Denmark (primary) and Sweden (secondary, since 2025-Q4). Roll-out to Norway is on the 2027 roadmap and is blocked pending a Norwegian-language evaluation set (§3.1).
- **Sectors:** medium-sized employers (50–500 FTE) in IT services, professional services, retail head-office functions and light manufacturing. Sold B2B per recruiter seat.
- **Scale (2026-H1):** 214 deployer tenants, 1 870 active recruiter seats, approximately 61 000 applications screened per month.

1.3 Version and release identification

- **Product version:** v2.4.0, released 2026-04-22; semantic versioning.
- **Model artefact:** `acmescreen-xlmr-da-sv-en-2.4.0`, SHA-256 `b7a1...c4e2`; the full digest is recorded in `legal/MODEL_DIGEST_2026-05-07.txt`.
- **Superseded version:** v2.3.4, decommissioned 2026-04-29. The artefact, its training manifest and its evaluation results are retained together with this technical documentation for **10 years** after the system is placed on the market, in accordance with **Art. 18** (documentation keeping). Earlier versions of this document cited Art. 12(1) for that retention; Art. 12 governs the *logging capability* of the system, not the retention of documentation, and the citation has been corrected (§12).

- **Update policy:** minor releases are covered by the pre-determined change plan in §12.2 and do not constitute a substantial modification within the meaning of Art. 3(23).

1.4 Hardware and software architecture

- **Inference runtime:** Python 3.12, PyTorch 2.4.1, HuggingFace transformers 4.46.2, sentence-transformers 3.1.1, served via TorchServe 0.12 behind FastAPI 0.115.
- **Orchestration:** Kubernetes 1.30 on AWS EKS in eu-central-1 (Frankfurt); node pool m6i.2xlarge for inference and g5.xlarge (NVIDIA A10G) for the embedding stage.
- **Storage:** PostgreSQL 16 (RDS, encrypted at rest with a customer-managed KMS key), S3 for model artefacts (versioned, MFA-delete), CloudWatch and Grafana Cloud for telemetry.
- **Network boundary:** VPC-private. The inference service is reachable only through the Acme application API gateway; there is no public ingress to inference pods.
- **Forms in which the system is placed on the market** (Annex IV point 1(d)): exclusively as a multi-tenant SaaS module within the Acme HR-tech platform, accessed through the platform UI and a REST API. There is no on-premises, embedded or downloadable distribution.
- **Third-country transfers:** none. Training and inference are performed in eu-central-1 only; no personal data is transferred outside the EEA. Sub-processors are listed in the DPA annex.

1.5 User interface provided to the deployer

Annex IV point 1(g).

The recruiter-facing interface is the control surface through which human oversight under Art. 14 is actually exercised, so it is documented here as part of the system rather than as a cosmetic layer.

Element	Behaviour	Oversight function
Ranked shortlist	The top N candidates for the posting, N configurable 10–50, default 20. Ranks 21–30 are shown collapsed beneath the fold and can be expanded at any time	Art. 14(4)(b), (d)
Confidence band	A banded indicator (low / medium / high) rather than a bare decimal, with the numeric score available on hover	Art. 14(4)(c)
"Why this rank" panel	The three structured features contributing most to the score, from a post-hoc SHAP attribution, each rendered as a plain-language statement	Art. 14(4)(c)
Model card panel	The validated envelope, known limitations and current subgroup metrics for the release in use, in-product and always reachable	Art. 14(4)(a)
Suitability banner	Warns the recruiter when the posting falls outside the validated envelope — a job family with fewer than 200 historical postings, or a high share of very short CVs	Art. 14(4)(a)
Reason-code modal	Required on every rejection; blocks the action where the selected reason maps to a protected characteristic	Art. 14(4)(d); GDPR Art. 22
Admin console	Tenant-wide stop control, N configuration, oversight-audit export, log retrieval	Art. 14(4)(e); Art. 26(6)

The interface is available in Danish, Swedish and English and is tested against WCAG 2.2 level AA; the most recent accessibility audit is dated 2026-03-19. Annotated screenshots of each element are held at docs/UI_REFERENCE_AcmeScreen_v2.4.0.pdf and form part of this technical documentation.

1.6 Instructions for use provided to the deployer

Annex IV point 1(h); Article 13.

The instructions for use are issued as a versioned, controlled document — `docs/IFU_AcmeScreen_v2.4.0.pdf`, in Danish, Swedish and English — delivered to every deployer before the system is put into service and re-issued with every minor release. Article 13(3) prescribes their content, and the mapping is as follows.

Art. 13(3)	Required content	Where it is provided
(a)	Identity and contact details of the provider	IFU §1; Acme HR-tech ApS, CVR 999999999, with a named product contact and the security contact
(b)	Characteristics, capabilities and limitations of performance, including accuracy, robustness and cybersecurity metrics; foreseeable risks; performance for specific persons or groups; input-data specifications	IFU §§3–5, reproducing §2.6, §3.1, §3.2, §3.3 and §3.4 of this file, including the open deviation recorded in §4.4
(c)	Changes to the system and its performance pre-determined at the time of the initial conformity assessment	IFU §6, reproducing the pre-determined change plan in §12.2
(d)	Human oversight measures, including the technical measures that facilitate interpretation of the output	IFU §7, reproducing §6
(e)	Computational and hardware resources, expected lifetime, and the maintenance and care measures needed, including the frequency of software updates	IFU §8. As a SaaS module the deployer needs no local compute; expected service life of the v2 line is to 2029-12-31, with quarterly model refreshes and monthly platform releases
(f)	Mechanisms allowing the deployer to collect, store and interpret the logs	IFU §9, describing the log export in the admin console and the deployer's own retention duty of at least six months under Art. 26(6)

The instructions for use also state expressly what the deployer must do under Art. 26 and what falls outside the system's validated envelope; that allocation is set out in §6.6.

Article 50(1) is not engaged, because AcmeScreen does not interact directly with natural persons — the applicant does not converse with the system. The obligation to inform applicants that a high-risk system is used in a decision affecting them arises instead under **Art. 26(11)**, and rests on the deployer; the product provides the notice text and a placement recommendation (§6.6).

§2 — Detailed description of elements and development process

Annex IV point 2. Cross-references: Art. 10 (data governance), Art. 15 (accuracy).

2.1 Training data

- **Source corpus:** **pseudonymised** CVs collected through Acme's own platform between 2020-01-01 and 2024-12-31 — n = 420 000 unique CVs across 11 800 job postings.
- **Pseudonymisation pipeline:** every element of the canonical direct-identifier list **PII-CANON-v3** (defined once in §5.4 and used identically throughout this document) is removed before the record reaches the training warehouse. Original identifiers remain only in the operational database, under separate access control. Reviewed by the DPO in November 2024; minute at `legal/DPIA_2026-05-07.md` §A.3.
- **Terminology.** This corpus is **pseudonymised, not anonymised**. Records remain re-identifiable by Acme through the retained tenant and submission keys, and free-text narrative can carry indirect identifiers. It is therefore personal data and the GDPR applies to it in full. Document version 1.0 described the corpus as "anonymised" while simultaneously relying on GDPR Art. 6(1)(f); those two statements cannot both be true, because genuinely anonymous data would require no lawful basis at all. The terminology has been corrected throughout (§5.6, §12).
- **Languages:** Danish approximately 63 %, Swedish 24 %, English 13 %.
- **Splits:** 80 / 10 / 10 train / validation / test. The test set was frozen on 2024-12-31 and has never been used in a hyperparameter sweep.
- **Labels:** a binary `passed_screen` flag derived from recorded recruiter actions — advanced to interview versus archived. De-confounded for recruiter identity and calendar week to reduce reviewer-specific effects.

2.2 Feature engineering

CV text is parsed by a deterministic Danish and Swedish NLP pipeline — spaCy 3.7 with `da_core_news_lg` and `sv_core_news_lg` — into structured fields: years of experience, education level mapped to ISCED, declared skills, language proficiencies and free-text narrative. The model consumes the narrative together with a structured-feature prompt.

No protected or special-category attribute is a feature of the model. Gender, ethnicity, religion, health, trade-union membership, sexual orientation and date of birth are removed before tokenisation; age is used only as the binary "meets the statutory minimum working age" flag required by Danish law. Employment-gap features were removed from the structured set in v2.2.0 because they act as a proxy for parental leave, and therefore for sex and age.

This is a design choice, not a legal defence. Excluding a protected attribute from the feature set does **not** prevent indirect discrimination through correlated features, which is precisely why the measurement described in §5.2 is carried out.

2.3 Model architecture and training

- **Architecture:** **XLM-RoBERTa base** — 12 transformer layers, 768 hidden units, 12 attention heads, SentencePiece vocabulary of 250 000, approximately 270 M parameters — initialised from the public `xlm-roberta-base` checkpoint and fine-tuned on the Acme training set with a binary classification head. A multilingual checkpoint was selected because the corpus spans Danish, Swedish and English, and a Danish-only or English-only encoder underperformed on the cross-lingual evaluation by 4 to 6 points of Recall@20. Document version 1.0 described the architecture as "BERT-base ... multilingual checkpoint `xlm-roberta-base`", which named two different architectures in a single sentence; the description and the artefact name have been corrected (§12).
- **Hyperparameters:** learning rate 2e-5, batch size 32, 3 epochs, warm-up 10 %, weight decay 0.01, AdamW.
- **Compute:** four A10G GPUs for approximately 14 hours per training run.
- **Reproducibility:** a locked `requirements.lock`, a pinned dataset content hash and a fixed random seed; each run is reproducible to within 0.004 F1.

- **Tooling:** Weights & Biases for experiment tracking (project `acmescreen-prod`) and MLflow 2.16 as model registry. Training runs, evaluation results and model cards are retained for **10 years** together with this file under **Art. 18**. Document version 1.0 cited Art. 12 for that retention; Art. 12 concerns the logging capability of the system, and the citation has been corrected (§12).

2.4 Third-party libraries

Library	Version	Licence	Purpose
torch	2.4.1	BSD-3	Tensor and autograd
transformers	4.46.2	Apache-2.0	Model loading
sentence-transformers	3.1.1	Apache-2.0	Embeddings
spacy	3.7.6	MIT	NLP parsing
fairlearn	0.11.0	MIT	Bias diagnostics
evidently	0.4.33	Apache-2.0	Drift detection
fastapi	0.115.0	MIT	Serving API

A full software bill of materials in CycloneDX format is maintained at `legal/SBOM_acmescreen_2.4.0.cdx.json` and renewed on every release, following ISO/IEC 5230 practice for open-source licence compliance.

2.5 Data labelling protocol

Recruiter "advance to interview" events are the ground-truth proxy. The labelling protocol at `legal/labelling_protocol_v2.pdf` requires recruiters to log a structured reason code when archiving a candidate; reason codes mapping to non-merit grounds, such as "position withdrawn", are excluded from training. Inter-rater reliability is sampled monthly — Cohen's $\kappa = 0.71$ on a 200-CV double-labelled subset, against a target of ≥ 0.65 .

The limitation of this design is stated plainly in §2.7: the labels record what recruiters *did*, not what a correct decision was.

2.6 Validation methodology and performance metrics

Measurement follows ISO/IEC TS 4213:2022 for the assessment of machine-learning classification performance. Results on the frozen held-out test set, $n = 42\,000$, with Wilson 95 % confidence intervals:

Metric	Value	95 % CI	Role
Precision@20	0.84	0.83 – 0.85	Acceptance gate — floor 0.80
Recall@20	0.79	0.78 – 0.80	Acceptance gate — floor 0.75
F1	0.81	0.80 – 0.82	Reported
AUC-ROC	0.88	0.87 – 0.89	Reported, not gated
Calibration error (ECE)	0.04	—	Acceptance gate — ceiling 0.06

This table is the single source of truth for the headline performance figures; §7.1 refers to it rather than restating the numbers. The subgroup breakdown is in §3.2.

The examination of the training, validation and test sets for possible biases is required by **Art. 10(2)(f)**, and the adoption of appropriate measures to detect, prevent and mitigate those biases is required by **Art. 10(2)(g)**. Both are relied on: the examination is described in §5.2 and the mitigation measures in §4.2, §4.4 and §5.2. Document version 1.0 cited only Art. 10(2)(g) for the examination, conflating the two; the citation has been corrected (§12).

2.7 Appropriateness of the performance metrics

Annex IV point 4.

Annex IV point 4 asks not what the metrics are, but why they are the right ones for this system. The reasoning is set out below, including the respects in which the metric suite is inadequate on its own.

Why rank-sensitive metrics. The output is a *ranked shortlist* of a fixed size, not a binary accept or reject decision. Metrics computed at the operating point the product actually uses — the top N — therefore describe real behaviour, whereas a threshold-free summary would not. N = 20 is the default configuration and the metrics are recomputed for any tenant that changes N.

Why each metric is in the suite.

- **Recall@20** is the primary harm-relevant metric. The individual detriment that matters here is a qualified applicant never appearing on the shortlist and therefore never being seen by a human. Recall@20 measures exactly that failure, and it is the metric on which subgroup parity is assessed.
- **Precision@20** measures the proportion of the shortlist that recruiters went on to advance. It protects the utility of the product; it is not a fundamental-rights metric.
- **F1** is reported as a balance of the two and is used for release-over-release comparison.
- **AUC-ROC** is reported for comparability across retrains but is **not** used as an acceptance gate, because it aggregates performance over the whole ranking including regions the product never shows to a recruiter.
- **Expected calibration error** is gated because the interface exposes a confidence band to the human overseer. A systematically over-confident score would mislead the very person the AI Act relies on to provide oversight, so calibration is a precondition of effective Art. 14 oversight rather than a purely statistical nicety.
- **Equal-opportunity difference** — the true-positive-rate gap between subgroups — is the primary fairness metric, for the same reason Recall@20 is primary: it measures unequal rates of *failure to shortlist a qualified candidate*.
- **Demographic parity difference** is reported as a secondary diagnostic and is deliberately **not** an acceptance gate, because applicant pools differ legitimately across job families and an enforced parity of selection rates could itself distort outcomes.

Why raw accuracy is not used. The class balance is 41 % positive to 59 % negative; a degenerate classifier would score misleadingly well on accuracy. It is not reported.

Statistical adequacy. Subgroup metrics are reported only where the subgroup cell contains at least 500 test records. Below that threshold the result is recorded as "insufficient sample" rather than as a precise-looking but unstable number. Confidence intervals are Wilson intervals; overlapping intervals are treated as no measured difference.

Where the suite is inadequate — stated openly. The labels are a proxy for recruiter behaviour, not an objective measure of job performance. Every metric above therefore measures *agreement with historical human decisions*, and a model that reproduced historical human bias perfectly would score well on all of them. This is the central limitation of the design, and it is why the metric suite is not relied upon alone: it is paired with the bias diagnostics under Art. 10(5) in §5.2, the risk controls in §4, and the human-oversight measures in §6. Two further limitations follow from it: the metrics cannot detect a candidate who never applied because of the system's reputation, and they cannot measure the quality of hires that were never made.

§3 — Monitoring, functioning and control

Annex IV point 3. Cross-references: Art. 12 (logging), Art. 14 (oversight), Art. 15 (accuracy).

3.1 Capabilities and limitations in performance

What the system does. For a given job posting, AcmeScreen scores each application by the predicted probability that a trained recruiter would advance that candidate to interview, and returns the ranked top N. It operates on Danish, Swedish and English free text together with the structured fields listed in §2.2.

Limitations, each of which is surfaced in the product and in the instructions for use:

Limitation	Measured effect	How it is surfaced
CVs shorter than 400 tokens	Precision@20 falls by 0.07	Per-posting suitability banner
Job families with fewer than 200 historical postings	Recall@20 falls by up to 0.11; confidence intervals widen	Suitability banner; model card
Norwegian-language CVs	Not validated — no evaluation set exists	Norwegian roll-out blocked at product level
Scanned or image-only CVs	No text layer; the system cannot rank them	Routed to a manual queue, never discarded
Cross-posting comparison	Scores are calibrated within a posting only	In-product notice; API returns a posting-scoped score
Credential verification	Not performed at all	Stated in the IFU and in the UI model card
Internal mobility, promotion, termination	Outside the validated envelope	Contractually excluded; blocked by the mandatory posting identifier

Overall expected level of accuracy in relation to the intended purpose is the table in §2.6. Deployers are told, in those terms, that roughly one in five candidates whom a recruiter would have advanced will not appear in a default shortlist of 20 — which is why the interface keeps ranks 21–30 one click away and why the escalation in §6.4 exists.

3.2 Expected accuracy for specific groups of persons

Annex IV point 3 requires the degrees of accuracy for specific persons or groups on which the system is intended to be used. The canonical subgroup table below is the single source for these figures; §5.2 and §7.1 refer to it.

Subgroup (attribute source)	Cell size	Re-call@20	TPR difference vs. reference	Status
All applicants (reference)	42 000	0.79	—	Baseline
Female (gender proxy)	20 160	0.78	0.03 (demographic parity 0.03)	Within tolerance
Male (gender proxy)	21 420	0.80	Reference	Within tolerance
Non-binary or unclear	420	—	Insufficient sample	Not reported — see §2.7
Age band 25–34 (graduation-year proxy)	17 220	0.81	Reference	Within tolerance
Age band 45+ (graduation-year proxy)	6 720	0.77	0.04	Within tolerance
Danish-first CV	26 460	0.80	Reference	Within tolerance
Swedish-first CV	10 080	0.79	0.01	Within tolerance
English-first CV	5 460	0.74	0.06 — exceeds the 0.05 gate	Deviation DEV-2026-014, §4.4

The English-first gap is the one open non-conformity in this file. It is not presented as compliant: it breaches the release gate in §4.2, and the documented deviation, its compensating controls and its remediation plan are set out in §4.4.

Attribute values are proxies, derived as described in §5.2, and carry their own error: the forename-based gender proxy agrees with self-declared gender on 94 % of a 2 000-record consented validation subset, and the graduation-year age proxy on 89 %. Proxy error is a further reason the figures are treated as indicative and the gate is applied conservatively.

3.3 Foreseeable unintended outcomes and sources of risk

Source of risk	Fundamental right engaged	Control
Indirect discrimination through the language proxy	Non-discrimination; Charter Art. 21	Measured (§3.2); open deviation and shortlist extension (§4.4)
Proxy discrimination through CV gaps correlating with parental leave	Non-discrimination; sex and age	Employment-gap features removed in v2.2.0; residual free-text exposure monitored
Feedback loop reinforcing historical recruiter preference	Non-discrimination	Label de-confounding (§2.1); override-rate band monitor (§6.5)
Automation bias — the recruiter defers to the ranking	Right to an effective remedy; GDPR Art. 22	§6.4 controls; dwell-time and override audits
Shortlist truncation — a qualified candidate is never seen	Access to employment	Recall@20 gate; ranks 21–30 one click away; periodic prompt
Adversarial CV optimisation favouring resourced applicants	Fairness of process	Keyword-stuffing detection (§7.2)
Service failure during a live campaign	—	Fails to chronological order; no candidate is hidden (§3.4)
Scope creep into promotion or termination decisions	Non-discrimination; Art. 22 GDPR	Contractual exclusion; posting-identifier enforcement

None of the practices prohibited by Art. 5 is engaged. In particular the system performs no emotion inference in the workplace within the meaning of Art. 5(1)(f), no biometric categorisation under Art. 5(1)(g), and no social scoring under Art. 5(1)(c). This screening was reviewed against each limb of Art. 5 on 2026-07-14; the assessment is at `legal/ART5_SCREENING_2026-07-14.md`.

3.4 Input-data specifications

Parameter	Specification	Behaviour on breach
Accepted formats	PDF with a text layer, DOCX, TXT	Image-only files are routed to a manual queue
Maximum size	20 MB per file	Rejected at upload with a message to the applicant
Maximum length	32 000 tokens	Truncated with a flag; the recruiter sees a truncation notice
Minimum length	400 tokens for a reliable score	Scored but flagged as low-confidence
Required metadata	External job-posting identifier and posting text	Request refused — this is what blocks out-of-scope use
Character encoding	UTF-8, normalised to NFC	Normalised silently
Language	Danish, Swedish or English, auto-detected per paragraph	Other languages flagged as outside the validated envelope

Fail-closed rule. Where an application cannot be processed for any reason, it is placed in the manual review queue and the recruiter is notified. The system never silently drops a candidate, and a scoring failure never produces a rejection.

3.5 Observability and alerting

- **Metrics:** Prometheus scrape every 15 seconds — request latency, error rate, model-confidence histogram, feature-distribution drift, per-tenant shortlist composition.
- **Dashboards:** three Grafana boards — *Inference Health*, *Distribution Drift* and *Bias Monitor*. All three are mirrored read-only to the DPO and to the deployer's own administrator, so that the deployer can discharge its monitoring duty under Art. 26(5).
- **Alerting:** PagerDuty for P1 (inference 5xx above 1 % over 10 minutes, or mean latency above 500 ms over 5 minutes); Slack `#ml-ops` for P2 (drift KS above 0.15); a weekly digest for P3 (gradual subgroup TPR movement).

3.6 Drift detection

A weekly batch job runs a two-sample **Kolmogorov–Smirnov** test on each numeric input feature and computes the **population stability index** for categorical ones, comparing the rolling seven-day inference window against the training reference distribution. An investigation ticket is raised where the KS p-value is below 0.01 *and* the effect size exceeds 0.10.

The model is never retrained automatically. Retraining is a documented decision taken by the ML lead together with the DPO under the change-management procedure in §12, precisely so that a distribution shift cannot silently alter the system's behaviour between conformity assessments.

3.7 Human control points in the processing flow

1. Every output carries a confidence band and a three-feature SHAP attribution (§1.5).
2. A candidate cannot move to `rejected` without an authenticated recruiter action carrying a reason code; the API returns HTTP 422 for any attempt to do so without a valid recruiter session.
3. Bulk rejection is disabled above 50 candidates, as friction against mass-reject patterns.
4. Reason codes mapping to a protected characteristic are hard-blocked and routed to a compliance review queue.
5. The tenant-wide stop control returns the posting to chronological ordering within 30 seconds.

These controls are described in full, and mapped to Art. 14(4), in §6.

3.8 Logging

Article 12; retention under Article 19; access under Article 21.

The system automatically records, over its lifetime, the events needed to identify situations that may present a risk under Art. 79(1) or a substantial modification, to support post-market monitoring under Art. 72, and to allow the deployer to monitor operation under Art. 26(5) — this being the content required by **Art. 12(1) and (2)**. Each inference event records: timestamp, system and model version, tenant and posting identifier, a hash of the input feature vector, the returned ranking, the confidence score, the subsequent recruiter action and the recruiter identifier.

- **Retention. Art. 19** requires the provider to keep the automatically generated logs for a period appropriate to the intended purpose and of **at least six months**. Acme retains them for **10 years** in cold storage with S3 Object Lock in governance mode, a longer period chosen to match Danish employment-law limitation periods and the Art. 18 documentation-retention window. Log integrity is signed daily with an AWS KMS HMAC.
- **Deployer duty.** The deployer must itself keep the logs under its control for **at least six months** under **Art. 26(6)**; the export mechanism is described in §1.6.
- **Access by authorities.** Logs and this documentation are provided to a national competent authority on **reasoned request** under **Art. 21**. Art. 12(3) is not the basis for that access — it prescribes minimum logging content for Annex III point 1(a) remote biometric identification systems only, and does not apply to this system. Document version 1.0 cited Art. 12(3); the citation has been corrected (§12).

§4 — Risk management system

Article 9; Annex IV point 5.

A continuous, iterative risk-management process runs across the entire lifecycle of the system, as required by **Art. 9(1) and (2)**, and is operated as part of the Acme information-security management system. It is reviewed **quarterly** and on every minor version. The extract below is drawn from the live register at `legal/RISK_REGISTER_2026-05-07.xlsx`.

4.1 Identified risks

1. **R-01 — Bias on protected characteristics.** Disparate true-positive rates across gender, age band, or national and ethnic origin, including through proxies such as the language of the CV.
2. **R-02 — Attributed qualifications.** The model surfacing a skill or qualification not present in the CV text.
3. **R-03 — GDPR Art. 22 over-reach.** The screening becoming, in practice, a decision based solely on automated processing without meaningful human involvement.
4. **R-04 — Security breach.** Exfiltration of applicant data, model weights or inference logs.
5. **R-05 — Data poisoning.** Adversarial CV submissions designed to degrade or skew the model at retraining.
6. **R-06 — Loss of human oversight effectiveness.** Formal controls present but not exercised — automation bias, rubber-stamped reason codes, oversight assigned to staff without authority to override.

4.2 Risk-mitigation table

ID	Risk	Inherent	Mitigation	Residual	Owner	Re-evaluation
R-01	Bias on protected attributes	High	Protected attributes excluded from the feature set; quarterly Fairlearn audit under Art. 10(5); release gate RG-03 — block release where any subgroup TPR difference exceeds 0.05; shortlist extension where the gate is breached	Medium	Head of ML	Quarterly
R-02	Attributed qualifications	Medium	Output is a ranking over the input CV, not generated text; SHAP traces verify that every surfaced skill exists in the source; nightly grounding test on a 500-CV regression set	Low	ML Lead	Monthly
R-03	GDPR Art. 22 over-reach	High	Mandatory authenticated human action before any rejection, enforced in UI and API; flow-down to deployers in DPA §7.3; quarterly audit that at least 99 % of rejections show 5 s or more of recruiter dwell time	Medium	DPO	Quarterly
R-04	Security breach	High	SOC 2 Type 1 in place; annual external penetration test and on every major version; KMS encryption at rest; least-privilege IAM; disclosure programme at <code>security@acme-hrtech.dk</code>	Low	CISO	Continuous
R-05	Data poisoning	Medium	Training data drawn only from authenticated employer tenants, never scraped; anomaly detection on submission velocity per source; manual approval gate for every retraining dataset	Low	Head of ML	Per re-train
R-06	Loss of oversight effectiveness	High	Override-rate band monitor 5–25 %; dwell-time audit; reason-code consistency sampling; mandatory recruiter training with annual refresh	Medium	DPO	Monthly

Residual risk is judged acceptable under Art. 9(5) for every item above **except** the open deviation recorded in §4.4, which is accepted only on the conditions stated there and only until 2026-11-30.

4.3 Continuous review

A quarterly Risk Council — CTO, DPO, Head of ML, CISO and external counsel — reviews the register, the post-market monitoring findings and every open deviation. Material changes feed §11. The most recent Risk Council was held on **2026-07-28**; the next is scheduled for **2026-10-27**. Incident response is governed by `legal/INCIDENT_RESPONSE_2026-05-07.md` and connects to the Art. 73 reporting duty described in §11.2.

4.4 Documented deviation and acceptance decision — DEV-2026-014

DEVIATION RECORD DEV-2026-014 — RELEASE GATE RG-03 EXCEEDED

Field	Entry
Control breached	RG-03 — no subgroup true-positive-rate difference above 0.05 (R-01, §4.2)
Observation	Recall@20 difference of 0.06 for English-first CVs against Danish-first CVs, measured 2026-04-18 on release candidate v2.4.0-rc3, frozen test set n = 42 000, English-first cell n = 5 460 (§3.2)
Scope assessment	Language of CV is not itself a protected characteristic under the Danish Act on the prohibition of discrimination in the labour market. It is treated as within the scope of RG-03 because it correlates with national and ethnic origin and can therefore give rise to indirect discrimination
Decision	Release approved subject to conditions. Residual risk accepted as <i>Medium</i> , time-boxed
Decided by	Risk Council, 2026-04-20 — CTO, DPO, Head of ML. Re-affirmed 2026-07-28
Expiry	2026-11-30 , or the release of v2.5.0, whichever is earlier

Reasons for the decision. Four considerations were weighed and recorded.

1. The measured gap is a *recall* gap on a shortlist. It does not by itself reject any candidate; rejection requires an explicit authenticated recruiter action (§3.7, §6.1).
2. The compensating controls below were measured to reduce the gap on the affected cohort to 0.02, which is inside the gate.
3. Blocking the release would have retained v2.3.4 in production, whose measured difference on the same subgroup was **0.09** — that is, materially worse for the same affected group. Blocking would have harmed the people the gate exists to protect.
4. The obvious alternative — rebalancing by upsampling English-first CVs — could not be validated before the release window without contaminating the frozen test set, and a rushed fix carried a higher risk of an unmeasured regression.

Compensating controls, in force from release.

- Automatic extension of the shortlist from N = 20 to N = 25 for any posting where 15 % or more of applications are English-first; measured residual recall difference on that cohort **0.02**.
- A banner to the recruiter on such postings, stating that the shortlist has been extended and why.
- Bias re-check on this subgroup moved from quarterly to **monthly** (§11.1).
- The deviation is disclosed to deployers in the instructions for use (§1.6), so that they can discharge their own Art. 26 duties with knowledge of it, and to applicants on request under Art. 86.

Remediation plan. Targeted upsampling of English-first CVs and a balanced Danish, Swedish and English fine-tuning set in v2.5.0; construction of a new frozen test set with a minimum English-first cell of 12 000; target difference ≤ 0.03 . Owner: Head of ML. Due **2026-11-30**. Tracked as R-01-D1 in the risk register and reported in the 2026-Q2 and 2026-Q3 post-market monitoring reports.

Escalation if the target is missed. The acceptance lapses automatically on expiry. RG-03 then blocks the v2.5.0 release, no further deviation may be granted at Head-of-ML or Risk Council level, and the decision escalates to the board of Acme HR-tech ApS with the DPO's written opinion attached. A gate that is quietly overridden is worse than no gate at all, which is why this breach is stated, bounded, compensated and given an expiry date rather than absorbed.

4.5 Criteria for judging residual risk acceptable

Article 9(5) requires that the residual risk associated with each hazard, and the overall residual risk, be judged **acceptable**. A judgement of that kind is only meaningful if the criteria are written down in advance, so they are recorded here and applied at every quarterly Risk Council.

A residual risk is judged acceptable only where all four of the following hold:

1. **No unmitigated pathway to an individual detriment.** No single failure of the system can, on its own, reject an applicant. Every rejection requires an authenticated human action (§3.7, §6.2).
2. **The risk is measured, not assumed.** A quantified indicator exists, is monitored at the cadence in §11.1, and has a defined threshold. A hazard that cannot yet be measured is not accepted; it is escalated.
3. **The residual level is at or below the declared tolerance,** or is covered by a deviation record with compensating controls, a named owner, a remediation plan and an expiry date (§4.4).
4. **The affected person retains a route to challenge the outcome** — through the deployer under Art. 26(11), the explanation right under Art. 86, and the GDPR Art. 22 safeguards.

The overall residual risk is judged acceptable where every individual residual risk meets those criteria **and** no two Medium residual risks share a single point of failure. On the last assessment, 2026-07-28, R-01, R-03 and R-06 were each Medium; they were confirmed not to share a common control, since they are mitigated respectively by the data pipeline, the application layer and the training and audit programme. The overall residual risk was accordingly judged acceptable, subject to the expiry of DEV-2026-014 on 2026-11-30.

Acceptance is recorded by name. It is signed jointly by the CTO and the DPO; neither can accept a residual risk alone, and neither may accept a risk rated High.

§5 — Data and data governance

Article 10; Annex IV point 2.

5.1 Training-set composition

- 420 000 CVs, deduplicated by content hash.
- Job categories, top five: IT and software 31 %, sales 14 %, administration 12 %, customer service 9 %, engineering 8 %.
- Time coverage 2020–2024. Class balance at training time, after stratified sampling: 41 % positive, 59 % negative.
- Relevance, representativeness and error rates are assessed against the intended purpose as required by **Art. 10(3)**; the corpus over-represents IT and software roles relative to the Danish labour market, which is recorded as a known limitation and is the reason for the job-family suitability banner in §1.5.

5.2 Bias examination and the legal basis for it

Article 10(5) is the provision that permits this processing. By way of derogation from Art. 9(1) GDPR, Art. 10(5) of the AI Act allows a provider of a high-risk AI system to process special categories of personal data **to the extent strictly necessary for the purpose of ensuring bias detection and correction**, subject to the safeguards it lists. Acme relies on Art. 10(5) as the lawful gateway for the diagnostics described here.

Document version 1.0 cited Art. 10(5) for the opposite proposition — as a reason *not* to process protected attributes. That inverted the provision. Art. 10(5) is an enabling derogation, not a prohibition, and the argument has been rewritten (§12).

Art. 10(5) condition	Implementation
(a) The purpose cannot be fulfilled by processing other data, including synthetic or anonymised data	Documented alternatives assessment of 2026-03-11: synthetic data reproduced the training distribution but could not measure the <i>observed</i> disparity in real recruiter behaviour, which is the quantity under audit. Held at <code>legal/ART10_5_NECESSITY_2026-03-11.md</code>
(b) Technical limitations on re-use, and state-of-the-art security and privacy-preserving measures	Diagnostics run in a segregated AWS account with no egress; attribute values are key-salted and survive only as aggregate counts
(c) Strict access control	A named list of four — Head of ML, DPO, one data engineer, the external auditor; hardware-token MFA; every access logged and reviewed monthly by the DPO
(d) No transmission or transfer to, or other access by, third parties	No transmission outside the diagnostic account; the external auditor works inside it under an Art. 28 processing agreement
(e) Deletion once the bias has been corrected or at the end of the retention period, whichever is first	Attribute-level data deleted at the close of each audit cycle, maximum 90 days; last deletion run 2026-08-03, evidenced in <code>dataops/art10_5_deletion_log.csv</code>
(f) Records of processing activities under GDPR Art. 30 stating why processing special categories was strictly necessary	ROPA entry ACME-ROPA-017, last reviewed 2026-06-30

Attribute sources. For the historical corpus, protected attributes were never collected, so proxies are used and are labelled as such: forename pattern for gender, graduation year for age band, and first-language metadata for a broad national-origin signal. Proxy accuracy is itself validated (§3.2). From 2026-Q3 a voluntary, separately stored self-declaration channel supplies actual attributes for a consented audit cohort, used only inside the segregated environment under Art. 10(5).

Measured distributions on the training data: gender proxy 48 % female, 51 % male, 1 % non-binary or unclear; age bands 18–24 at 19 %, 25–34 at 41 %, 35–44 at 24 %, 45–54 at 11 %, 55+ at 5 %; CV language Danish-first 63 %, Swedish-first 24 %, English-first 13 %. The resulting subgroup performance is in the canonical table at §3.2, and the one breach of tolerance is handled in §4.4.

The distinction that matters. Two separate things are true, and version 1.0 confused them. First, the live inference path receives no protected or special-category attribute — a design choice (§2.2). Second, protected attributes *are* processed, in a controlled audit environment, in order to measure disparate impact — and Art. 10(5) is what makes that lawful. Excluding an attribute from the model is not a defence against indirect discrimination; it merely makes discrimination harder to see. Measuring it is the defence, and the AI Act expressly provides for it.

5.3 Data-quality controls

Pre-ingestion: schema validation, encoding normalisation to UTF-8 NFC, language detection, duplicate detection and rejection of malformed files. Post-ingestion: a data-quality analyst spot-checks 100 randomly selected CVs each week, with results logged in `dataops/qc_log.md`. Data-quality criteria and the examination for errors are those required by Art. 10(2)(c)–(e).

5.4 Data-cleaning protocol and the canonical identifier list

PII-CANON-v3 — canonical direct-identifier list. This list is defined here once and is the same list referred to in §2.1; the two sections previously disagreed, and that inconsistency has been resolved (§12).

name; postal address; email address; telephone number; photograph and photographic EXIF metadata; **date of birth**; CPR-nummer or personnummer; passport or national identity number; **IBAN or other bank account identifier**; driving-licence number; personal website and social-media URLs.

Removal is *additive-deny*: a union of regular-expression and named-entity-recognition detectors strips every element of PII-CANON-v3. Any CV in which more than five identifier tokens remain flagged after the first pass is quarantined and reviewed by a human before it is either deleted or admitted to the corpus. The same list governs the inference path, where identifiers are stripped before tokenisation.

5.5 Data protection impact assessment

A DPIA under GDPR Art. 35 was completed on 2024-11-18 and refreshed on 2026-04-30; it is held at `legal/DPIA_2026-05-07.md`. It covers necessity and proportionality, risks to the rights and freedoms of applicants, the safeguards described in this file, and the conclusion that prior consultation with Datatilsynet under Art. 36 GDPR was not required because the residual risk after mitigation is not high. That conclusion is revisited at each quarterly review and would be reopened by any material change to §4.4.

5.6 Lawful basis under the GDPR

Roles. For the screening service Acme acts as **processor** within the meaning of GDPR Art. 4(8); each employer-customer is the **controller**. A processor holds no lawful basis of its own for the controller's processing — it processes only on documented instructions under Art. 28(3)(a). Document version 1.0 stated the basis as though Acme held it directly; that was wrong in principle and has been rewritten (§12).

The lawful basis is therefore the controller's, and the DPA records which basis the controller has designated:

- **Art. 6(1)(b), second limb** — processing necessary in order to take **steps at the request of the data subject prior to entering into a contract**. This is the applicable limb in a recruitment context: no contract yet exists between applicant and employer, and it is the applicant who initiates the application. The first limb, performance of a contract to which the data subject is already party, does **not** apply and is not relied on. Version 1.0 relied on the first limb; the citation has been corrected.

- **Art. 6(1)(f)** — where a controller instead relies on legitimate interests in efficient candidate triage, the balancing test is the controller's own responsibility. Acme supplies a template at `legal/LIA_TEMPLATE_2026-08-31.md`; it does not complete it on the controller's behalf.

For the training corpus Acme is a separate **controller** and relies on **Art. 6(1)(f)**, legitimate interest in product improvement. Because the corpus is **pseudonymised and not anonymous** (§2.1), it remains personal data and the GDPR applies in full — which is what makes a lawful basis necessary in the first place. The balancing test is at `legal/LIA_2026-05-07.md`. Applicants are informed at submission and may opt out; the opt-out propagates to the training warehouse within 30 days. Tenants whose own terms are incompatible with the training use are excluded from the corpus.

Applicant rights. Access, rectification, erasure and objection are served through the deployer as controller, with Acme supporting under Art. 28(3)(e). Rights requests concerning the training corpus are served by Acme directly. The right to obtain an explanation of an individual decision under **Art. 86** of the AI Act is supported by the per-candidate attribution described in §1.5.

5.7 Special-category data

Special categories under GDPR Art. 9(1) — health, religion or belief, trade-union membership, racial or ethnic origin, sexual orientation, biometric and genetic data — are **blocked from the production path** by a named-entity and keyword denylist applied both to the training corpus and to the live inference path. Any CV whose special-category content cannot be redacted with high confidence is excluded from the training corpus.

No GDPR Art. 9(2) condition is required for the production path, because no special-category data is processed there. For the **bias-audit path** described in §5.2, special-category data is processed under the **Art. 10(5) AI Act derogation**, which operates by way of derogation from Art. 9(1) GDPR and carries its own conditions, implemented as tabulated above. Version 1.0 stated flatly that no Art. 9(2) condition was relied on, without distinguishing the two paths; the position has been clarified (§12).

§6 — Human oversight

Article 14; Annex IV point 3. Referred to from §1.1, §3.7, §4.2 and §11.3.

6.1 Oversight design and the allocation between provider and deployer

The product is built on the principle that the system augments the recruiter and never replaces her. **Art. 14(3)** requires the provider to identify oversight measures and to state which are built into the system and which the deployer must implement. That split is set out expressly:

Measure	Built in by the provider — Art. 14(3)(a)	To be implemented by the deployer — Art. 14(3)(b)
Authenticated human action required before any rejection	Yes, enforced in UI and API	Must not attempt to automate the recruiter session
Reason code on every rejection	Yes	Must configure its own reason-code taxonomy from the supplied base set
Hard block on protected-characteristic reason codes	Yes	Must staff the resulting compliance review queue
Stop control returning the posting to chronological order	Yes	Must nominate administrators authorised to use it
Model card, limitations and suitability banners	Yes	Must ensure recruiters are trained to read them
Competence, training and authority of the overseer	Training material supplied	Deployer duty under Art. 26(2) — must assign oversight to a person with the necessary competence, training and authority
Information to workers' representatives and affected workers	Notice text supplied	Deployer duty under Art. 26(7)
Information to the applicant that a high-risk system is used	Notice text and placement guidance supplied	Deployer duty under Art. 26(11)

6.2 Measures giving effect to Article 14(4)

Art. 14(4)	Requirement	Implementation
(a)	Understand capacities and limitations; monitor operation; detect anomalies and unexpected performance	In-product model card carrying the §3.1 limitations and the current §3.2 subgroup figures; per-posting suitability banner; anomaly alerts routed to the tenant administrator as well as to Acme
(b)	Remain aware of automation bias	No auto-reject capability exists anywhere in the product; ranks 21–30 are always one click away and a prompt to review them appears on every fifth posting; confidence is shown as a band, not a false-precision decimal; the override-rate monitor in §6.5 detects crews that have stopped exercising judgement
(c)	Correctly interpret the output	Three-feature SHAP attribution rendered in plain language; an in-UI statement that scores are comparable within a posting only; the published calibration curve; tooltips explaining what a given confidence does and does not mean
(d)	Decide not to use the system, or to disregard, override or reverse the output	Any candidate at any rank can be advanced in one action; a rejection can be reversed for 30 days; every override is logged and feeds §6.5. Friction is deliberately placed on rejection and never on inclusion — advancing a low-ranked candidate requires no justification
(e)	Intervene or interrupt through a stop button	Admin-console stop control disables ranking tenant-wide within 30 seconds and falls back to chronological ordering; a global control is available to Acme SRE; tested monthly, last test 2026-08-04 at 12 seconds observed

On the safe state. The fallback deliberately shows *every* applicant in chronological order rather than showing none. A stop control that hid candidates would convert an availability failure into a fundamental-rights failure.

Art. 14(5), which requires separate confirmation by two natural persons, applies to Annex III point 1(a) remote biometric identification systems and does not apply to this system.

6.3 Competence and training of the overseers

Every recruiter completes an eight-hour module before activation: Module A on how the system works and what it does not do; Module B on bias awareness, protected characteristics and applicant rights under the GDPR; Module C hands-on with the override and reason-code interface; Module D on spotting and reporting incidents. Completion is recorded and the certificate is refreshed annually; the log is at `Legal/TRAINING_LOG_2026-05-07.csv`.

This training also serves the **Art. 4** AI-literacy obligation, which has applied since **2 February 2025** and is therefore in force today, independently of the 2027-12-02 date for the Annex III obligations. Art. 4 binds both Acme and each deployer in respect of its own staff; the training material is supplied to deployers for that purpose, and the deployer remains responsible for discharging its own Art. 4 duty.

6.4 Automation-bias controls

Automation bias is treated as a distinct risk (R-06) rather than as a training topic, because it is the failure mode that would hollow out every other control in this section. Three controls operate: the absence of any bulk or automated rejection path; the periodic prompt to review beyond the shortlist; and the override-rate band described next, which measures whether oversight is actually being exercised rather than merely being available.

6.5 Monitoring the effectiveness of oversight

A monthly audit samples 200 rejections per deployer tenant and tests: dwell time before rejection; reason-code consistency; and the recruiter override rate, being the share of cases in which a recruiter advances a candidate ranked below the shortlist.

- Dwell time of 5 seconds or more on at least 99 % of rejections.
- Override rate within a **5–25 %** band. Below 5 % suggests the ranking is being trusted uncritically; above 25 % suggests it is not trusted at all and the model is not fit for that tenant's postings.
- Readings outside either band are escalated to the Risk Council and, where they concern one tenant, raised with that deployer under Art. 26(5).

Results for 2026-Q2: mean dwell 11.4 s, 99.4 % of rejections above threshold; override rate 14.2 % across all tenants, with two tenants below 5 % which are under active follow-up.

6.6 What the deployer must do

The following are **deployer** obligations under Art. 26 and apply from **2 December 2027**. They are reproduced in the instructions for use so that no deployer can be unaware of them.

Provision	Deployer obligation
Art. 26(1)	Use the system in accordance with the instructions for use
Art. 26(2)	Assign human oversight to natural persons with the necessary competence, training and authority, and with the support to exercise it
Art. 26(4)	Ensure that input data is relevant and sufficiently representative in view of the intended purpose
Art. 26(5)	Monitor operation, and inform the provider and the market-surveillance authority where a risk under Art. 79(1) is identified, suspending use where appropriate
Art. 26(6)	Keep the logs under its control for at least six months
Art. 26(7)	Before putting the system into service at the workplace, inform workers' representatives and the affected workers that they will be subject to it

Provision	Deployer obligation
Art. 26(11)	Inform the natural persons subject to its use that they are subject to a high-risk AI system

On the fundamental-rights impact assessment. Art. 27 requires a fundamental-rights impact assessment from deployers that are bodies governed by public law, private entities providing public services, or deployers of the systems in Annex III points 5(b) and (c). A private-sector employer deploying this employment-screening system is therefore **not** generally required to carry one out. Public-sector deployers and private operators providing public services **are**, and Acme supplies a data pack to support it. The Art. 26(7) duty to inform workers' representatives applies to every deployer regardless.

§7 — Accuracy, robustness and cybersecurity

Article 15.

7.1 Accuracy

The declared level of accuracy and the relevant accuracy metrics are those in §2.6, and the accuracy for specific groups is that in §3.2. They are stated once and not restated here, so that the two cannot drift apart. Both are declared to deployers in the instructions for use, as required by Art. 15(3).

Re-evaluation is quarterly, on a rolling holdout drawn from production traffic where the applicant has consented to evaluation use. A drop of more than 0.03 F1 against the release baseline triggers the investigation path in §11.1.

7.2 Robustness and adversarial testing

Three attack vectors were tested in 2026-Q1; the report is at `security/RED_TEAM_REPORT_2026-Q1.pdf` and the next red-team window is 2026-Q4.

1. **Keyword stuffing** — CVs padded with high-frequency positive tokens. Detection rate 96 %, using a length-normalised TF-IDF anomaly score.
2. **Paraphrase attack** — identical content rewritten to game tokenisation. Ranking variance below 5 percentile points, judged acceptable.
3. **White-text injection** — zero-opacity tokens embedded in a PDF. Caught in 100 % of cases by the text-layer hash check at ingest.

Resilience to faults and inconsistencies under Art. 15(4) is provided by the fail-closed rule in §3.4 and the chronological fallback in §6.2(e). The system has no continuous-learning path in production, so the feedback-loop risk that Art. 15(4) addresses for self-learning systems does not arise; retraining is a discrete, gated event under §3.6 and §12.

7.3 Cybersecurity

- **Threat model:** STRIDE, last refreshed 2026-03. Abuse cases include credential stuffing, model extraction, training-data exfiltration and prompt injection through CV content.
- **Penetration testing:** annual external, next scheduled 2026-09, plus internal quarterly.
- **Certification:** SOC 2 Type 1 issued 2026-02; Type 2 expected 2026-Q4. ISO/IEC 27001:2022 certified since 2024 (§9.2).
- **Vulnerability disclosure:** `security@acme-hrtech.dk`, 90-day policy, published at `acme-hrtech.dk/.well-known/security.txt`.
- **Patching:** critical CVEs within 7 days, high within 30, medium within 90.
- **Model-specific controls:** rate limiting and query-pattern analysis against model-extraction attempts; model weights encrypted at rest and never served to clients.

7.4 Failure modes

Maintained at `legal/FAILURE_MODES_2026-05-07.md`; each entry records the failure mode, its detection mechanism, the recovery procedure and the date last observed.

ID	Failure mode	Detection	Mitigation
F-01	Tokeniser out-of-memory on very long CVs	Pod OOM alert	32 000-token cap at ingest (§3.4)
F-02	Language detector misclassifies bilingual CVs as English	Weekly language-mix report	Per-paragraph detection since v2.3.0
F-03	Confidence calibration drift after a retrain	ECE gate in CI	Post-hoc Platt re-fit, gated before release

§8 — Risk management summary

This restates §4 in compact form for board and customer review. It is a summary; §4 governs in the event of any difference.

ID	Risk	Sever-ity	Likelihood	Mitigation summary	Resid-ual	Cadence
R-01	Disparate-impact bias	High	Medium	Feature exclusion, Art. 10(5) audit, RG-03 gate	Medium	Quarterly
R-02	Attributed qualifications	Medium	Low	Ranking-only design, SHAP grounding test	Low	Monthly
R-03	GDPR Art. 22 over-reach	High	Medium	Hard human-action gate, dwell-time audit	Medium	Quarterly
R-04	Security breach	High	Low	SOC 2, penetration testing, KMS, IAM	Low	Continu-ous
R-05	Data poisoning	Medium	Low	Authenticated sources, retrain gate	Low	Per retrain
R-06	Loss of oversight effectiveness	High	Medium	Override-rate band, dwell audit, training	Medium	Monthly

Residual-risk acceptance is signed by the CTO and the DPO at each quarterly Risk Council. Last sign-off **2026-07-28**. One deviation is open: **DEV-2026-014** (§4.4), expiring 2026-11-30.

§9 — Harmonised standards and other specifications applied

Annex IV point 7; Articles 40 and 41.

9.1 Status of harmonised standards

No harmonised standard has yet been cited in the *Official Journal* under **Art. 40** in respect of the Chapter III Section 2 requirements applicable to this system. The CEN-CENELEC JTC 21 work programme is under way but incomplete. The Commission has adopted no common specification under **Art. 41** for these requirements.

The consequence is stated plainly because it is often glossed over: **no presumption of conformity is available to this system today**. Conformity is therefore demonstrated by direct argument against the text of Arts. 9 to 15, which is what this document does. The standards below are applied voluntarily, as good practice and as a structuring device; applying them does not by itself establish conformity with the AI Act.

9.2 Standards applied

Standard	Edition	Role in this file	Clauses relied on	Evidence
ISO/IEC 42001	2023	AI management system framing the whole lifecycle	4, 6, 8, 9, 10	AIMS manual; internal audit 2026-06; certification targeted 2027
ISO/IEC 23894	2023	AI risk-management guidance mapped onto §4	Whole	Each register entry carries a 23894 control identifier
ISO/IEC 27001 + Annex A	2022	Information security	A.5.7, A.5.23, A.8.16	Certified since 2024; Statement of Applicability covers AcmeScreen; surveillance audit 2026-05
ISO/IEC 27701	2019	Privacy information management, GDPR interface	6, 7, 8	Extension of the 27001 scope
ISO/IEC 25059	2023	Quality model for AI systems; structures the metric selection in §2.7	5, 6	Metric rationale record
ISO/IEC TS 4213	2022	Measurement protocol for classification performance (§2.6, §3.2)	Whole	Evaluation protocol document
NIST AI RMF	1.0	Cross-check only; not a route to EU conformity	Map, Measure, Manage	Gap review 2026-02
prEN 18286	Draft	Tracked; not relied upon	—	Watch item; still draft as of 2026-08

Only clause references are given. The standards are copyright works available under licence and their text is not reproduced here.

9.3 What the standards do not cover

Applying ISO/IEC 42001 is not equivalent to complying with the AI Act, and a reader should not treat the table above as if it were. The gaps that matter:

AI Act requirement	Covered by a standard?	What actually discharges it here
Art. 9 risk management	Partly — ISO/IEC 23894 gives method, not the Art. 9(5) acceptability criterion	§4, including the acceptance decision in §4.4
Art. 10 data governance	Partly — no standard addresses the Art. 10(5) derogation	§5.2, §5.4

AI Act requirement	Covered by a standard?	What actually discharges it here
Art. 11 and Annex IV	No	This document
Art. 12 logging	Partly — 27001 A.8.15 covers logging generally, not Art. 12(2) content	§3.8
Art. 13 instructions for use	No	§1.6 and the IFU document
Art. 14 human oversight	No standard covers Art. 14(4) point by point	§6.2
Art. 15 accuracy and robustness	Partly — TS 4213 covers measurement, not the declared level	§2.6, §7

9.4 Update trigger

Where a harmonised standard giving a presumption of conformity is cited in the *Official Journal* under Art. 40, this file and the declaration in §10 are reviewed within 90 days and re-issued where the standard changes the basis of conformity.

§10 — EU declaration of conformity

Article 47; Annex V; Annex IV point 8.

The declaration below is drawn up under **Art. 47** and follows the content required by **Annex V**. It is signed on completion of the conformity assessment based on **internal control (Annex VI)** under **Art. 43(2)**. No notified body is involved: for an Annex III point 4 employment system the provider carries out the assessment itself, and the CE marking is then affixed under **Art. 48**.

EU DECLARATION OF CONFORMITY

Regulation (EU) 2024/1689, Article 47 and Annex V

TEMPLATE — NOT ISSUED. This declaration is unsigned and has no legal effect. It is completed and signed only when the internal-control conformity assessment under Art. 43(2) is complete.

Annex V	Field	Entry
	1 AI system name, type and reference	AcmeScreen — automated CV-screening high-risk AI system, v2.4.0. Model artefact acmescreen-xlmr-da-sv-en-2.4.0, SHA-256 b7a1...c4e2
	2 Name and address of the provider	Acme HR-tech ApS, [street], [postcode] [city], Denmark. CVR 99999999
	3 Sole responsibility	This declaration is issued under the sole responsibility of the provider
	4 Conformity statement	The AI system described above is in conformity with Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744
	5 Personal-data statement	The system processes personal data and complies with Regulation (EU) 2016/679
	6 Harmonised standards or common specifications	None available under Art. 40 or Art. 41 at the date of issue. Standards applied voluntarily are listed in §9.2
	7 Notified body	Not applicable. Conformity assessment is based on internal control under Annex VI, pursuant to Art. 43(2); no notified body is involved
	8 Place and date of issue; signature	To be completed on signature. Technical documentation reference: EU-DOC-AcmeScreen-v2.4.0

Signature — for and on behalf of Acme HR-tech ApS

Name and function: [name], Chief Executive Officer

Place and date of issue

The signed declaration is kept at the disposal of the national competent authorities for **10 years** after the system is placed on the market, under **Art. 47(2)**, and a copy is submitted to the relevant national competent authority on request. The same 10-year period applies to this technical documentation under **Art. 18**.

§11 — Post-market monitoring

Article 72; Annex IV point 9.

A post-market monitoring plan is maintained under **Art. 72(1)**. The Commission's implementing act laying down the template under Art. 72(3) has not yet been adopted; until it is, the plan follows the structure below and will be re-cast onto the template when it is published.

11.1 Monitoring plan

Cadence	Activity	Trigger threshold	Owner
Continuous	Latency, error rate, confidence histogram	5xx above 1 % over 10 min; mean latency above 500 ms over 5 min	SRE
Weekly	KS and PSI feature-drift batch	KS $p < 0.01$ and effect size > 0.10	ML Lead
Monthly	Bias re-check on the language subgroup	Any TPR difference above 0.05 — heightened cadence under DEV-2026-014	Head of ML
Monthly	Bias re-check, all other subgroups	Any TPR difference above 0.05	Head of ML
Monthly	Oversight-effectiveness audit (§6.5)	Dwell below 5 s, or override rate outside 5–25 %	DPO
Quarterly	Accuracy re-evaluation on consented holdout	F1 drop above 0.03 against baseline	Head of ML
Annually	External audit of the AIMS and the technical file	Any non-conformity	DPO and CTO

Findings are aggregated into a quarterly post-market monitoring report and archived alongside this file.

11.2 Reporting of serious incidents

Article 73 sets three distinct deadlines. They are frequently conflated, so they are set out separately here.

Provision	Situation	Deadline
Art. 73(2)	General rule for any serious incident	Immediately after the provider has established a causal link, or the reasonable likelihood of one, between the AI system and the incident — and in any event not later than 15 days after the provider, or where applicable the deployer, becomes aware of it
Art. 73(3)	A widespread infringement, or a serious incident consisting of a serious and irreversible disruption to the management or operation of critical infrastructure	Immediately, and not later than 2 days after becoming aware
Art. 73(4)	Where the serious incident has resulted in the death of a person	Immediately after the causal link is established or as soon as it is suspected, and not later than 10 days after becoming aware

Where necessary to ensure timely reporting, an initial incomplete report may be submitted and completed afterwards (**Art. 73(5)**). Document version 1.0 attributed the 10-day rule to Art. 73(2); it belongs to Art. 73(4), and the citation has been corrected (§12).

The 72-hour rule is not an AI Act deadline. Seventy-two hours is the GDPR **Art. 33** deadline for notifying a personal-data breach to Datatilsynet. NIS2 **Art. 23** operates a separate structure for in-scope entities — early warning within 24 hours, notification within 72 hours and a final report within one month. The three regimes are cumulative and independent; a single event can trigger all three, with different addressees and different clocks.

Addressees in Denmark. The market-surveillance authority for Annex III high-risk systems has not yet been designated by executive order (§R.3). Pending designation, **Digitaliseringsstyrelsen** is the addressee for AI Act notifications in its capacity as national contact point under lov nr. 467 af 14. maj 2025. **Datatilsynet** is the addressee for GDPR Art. 33 notification. Acme maintains pre-drafted notification templates for both.

11.3 Escalation thresholds

- **L1 — internal investigation.** Any single bias-gate breach, or any P2 alert.
- **L2 — DPO and CTO notified within 4 hours.** Repeated L1 in the same week, any P1 alert, or any applicant complaint alleging solely automated decision-making.
- **L3 — authority notification path opens.** Confirmed disparate-impact harm, a confirmed personal-data breach, a confirmed security incident exposing applicant data, or any event meeting Art. 73(1).

Each level has a named on-call owner; the rota is at `ops/ONCALL_2026-05.md`.

11.4 Sanctions exposure

The administrative fines that would apply to a failure of the provider obligations described in this file are those in **Art. 99(4)** — EUR 15 000 000 or 3 % of total worldwide annual turnover, whichever is higher. The full scale, including Art. 99(3) and Art. 99(5), is set out in **§R.4**. Acme's management is briefed on this exposure annually, and the briefing is minuted.

§12 — Record of changes through the lifecycle

Annex IV point 6.

12.1 Change record

Date	Item	Change	Substantial modification?
2026-02-11	v2.3.0	Per-paragraph language detection (F-02)	No — performance fix within the declared envelope
2026-04-22	v2.4.0	Retrain on data to 2024-12-31; shortlist extension control added under DEV-2026-014	No — pre-determined retrain, §12.2
2026-04-29	v2.3.4	Decommissioned; artefact retained under Art. 18	n/a
2026-08-31	Document v2.0	Documentation correction round — see §12.3	No — no change to the system

12.2 Pre-determined changes

Under **Art. 43(4)**, changes to a high-risk AI system that have been pre-determined by the provider and assessed at the time of the initial conformity assessment do not constitute a substantial modification. The following are pre-determined for this system and are declared in the instructions for use:

- Quarterly retraining on newer data from the same authenticated sources, with an unchanged architecture, an unchanged feature set and all §2.6 acceptance gates re-applied to a frozen test set.
- Recalibration of the confidence scale by post-hoc Platt scaling.
- Configuration of N within the range 10 to 50 per tenant.

Anything else — a change of architecture, the addition of a feature, an extension to a new language, or use for a purpose other than initial screening — is treated as a **substantial modification** within the meaning of Art. 3(23), and triggers a fresh conformity assessment and a new declaration.

12.3 Corrections made in document version 2.0

This document is itself under change control, and the corrections made in this version are recorded rather than made silently.

Location	Version 1.0	Version 2.0
Cover, §R.1	"per Art. 11(2)"	Art. 11(1) — 11(2) governs Annex I section A products
§1.3	Retention of a decommissioned version "per Art. 12(1)"	Art. 18 — documentation keeping
§2.3	Runs retained "≥10 years per Art. 12"	Art. 18
§2.6	Subgroup analysis "required by Art. 10(2)(g)"	Art. 10(2)(f) for the examination and 10(2)(g) for the mitigation
§3.8	Availability to authorities "per Art. 12(3)"	Art. 21 ; Art. 12(3) applies only to Annex III point 1(a)
§5.2	Art. 10(5) cited as a reason <i>not</i> to process protected attributes	Art. 10(5) restated correctly as the derogation permitting such processing for bias detection, with its conditions implemented
§11.2	10-day rule attributed to Art. 73(2)	Art. 73(4) ; 73(2), (3) and (4) now cited separately
§2.1, §5.6	Training corpus "anonymised" while relying on GDPR Art. 6(1)(f)	Pseudonymised throughout; the contradiction removed
§5.6	GDPR Art. 6(1)(b) first limb, held by Acme as processor	Second limb (pre-contractual steps), held by the controller ; processor role corrected
§2.1, §5.4	Two disagreeing identifier lists	Single canonical list PII-CANON-v3

Location	Version 1.0	Version 2.0
§2.3	"BERT-base ... xlm-roberta-base"	XLM-RoBERTa base only; artefact renamed <code>acmescreen-xlmr-da-sv-en-2.4.0</code>
§4.4	Bias gate breached but not addressed	Documented deviation DEV-2026-014 with reasons, compensating controls, remediation plan and expiry
§R.3	Danish authority described as future	Digitaliseringsstyrelsen , designated by lov nr. 467 af 14. maj 2025, in force 2025-08-02
§R.2	No applicability timeline	Timeline added, including Reg. (EU) 2026/1744
Throughout	A module name that does not correspond to any product PowerQuant sells; the domain <code>acmehr.dk</code> ; artefact paths under <code>canonical/legal/</code>	Modul 1 (fixed fee 10.999 kr); <code>acme-hrtech.dk</code> ; <code>legal/</code>

Annex A — Annex IV conformity map

Annex IV	Requirement	Where
1(a)	Intended purpose, provider, version	§1.1, §1.2, §1.3
1(b)	Interaction with hardware or other software	§1.4
1(c)	Software versions and update requirements	§1.3, §1.4, §12.2
1(d)	Forms in which the system is placed on the market	§1.4
1(e)	Hardware on which the system runs	§1.4
1(f)	Photographs or illustrations of a product component	Not applicable — SaaS only
1(g)	Basic description of the user interface for the deployer	§1.5
1(h)	Instructions for use for the deployer	§1.6
2	Elements and process of development	§2
3	Monitoring, functioning and control	§3
4	Appropriateness of the performance metrics	§2.7
5	Risk management system under Art. 9	§4
6	Relevant changes through the lifecycle	§12
7	Harmonised standards applied	§9
8	Copy of the EU declaration of conformity	§10
9	Post-market performance evaluation system under Art. 72	§11

Annex B — How PowerQuant produced this report

This document is a deliverable, not a generic template filled in from a precedent bank. It is grounded in artefacts produced during the engagement.

- **Modul 1 — AI inventory, Article 4 register and gap analysis.** Fixed fee **10.999 kr**, five working days. Produced the AI-system inventory, the Art. 4 AI-literacy register, the role-based control note, and the Annex III screening that classified AcmeScreen under **Annex III point 4(a)**. Every claim is source-referenced to the regulation itself. That classification is the foundation of this file: it determines that Arts. 9–15 apply, that the conformity route is internal control under Annex VI, and that the applicable date is 2 December 2027.
- **Annex IV depth.** The technical file itself is built on the Modul 1 output and maintained under an ongoing monitoring tier, where the Annex IV depth is set to the customer's needs. Scope and current pricing for every module are published at powerquant.dk/priser.
- **Companion artefacts** referenced in this file: `legal/DPIA_2026-05-07.md`, `legal/LIA_2026-05-07.md`, `legal/RISK_REGISTER_2026-05-07.xlsx`, `legal/INCIDENT_RESPONSE_2026-05-07.md`, `legal/ART10_5_NECESSITY_2026-03-11.md`, `legal/ART5_SCREENING_2026-07-14.md`, `security/RED_TEAM_REPORT_2026-Q1.pdf`.
- **How it is produced.** AI models draft; every regulatory claim is cross-checked against the text of the regulation; a named human reviewer approves each deliverable before it is released. PowerQuant is itself transparent about being an AI-based product under Art. 50.
- **What this is not.** It is not legal advice, and **PowerQuant ApS is not a law firm**. It is technical documentation meeting the formal requirements of Annex IV. Final legal interpretation belongs with the customer's own specialist IT or AI counsel. Conformity assessment, even by internal control under Annex VI, remains the provider's own responsibility.
- **Maintenance.** This file is versioned alongside the software. A minor release triggers a refresh of §2 and §7; a substantial modification under Art. 3(23) triggers a full re-issue and a fresh declaration under Art. 47.

PowerQuant ApS · CVR 46274067 · København, Danmark

kontakt@powerquant.dk · DPO: dpo@powerquant.dk · powerquant.dk

Document version 2.0 of 2026-08-31 · sample deliverable · fictional demonstration customer · not legal advice.

End of Annex IV technical documentation — AcmeScreen v2.4.0 — document version 2.0 of 2026-08-31.