Embed AI

AI ACT READINESS REPORT · 2026
Version 2026.09
2026-09-10 · Practical guide

ANNEX III POINT 4

AI in work and workforce management

Recruitment, selection and workforce management, from obligation to a review-ready evidence pack.

Recruitment | Workforce management | Bias | Human oversight

Legal reference date: 2026-09-10 · Practical guide, not formal legal advice.

FOREWORD FROM THE EXPERT

AI in work and workforce management.

Employers, recruiters and HR-tech vendors increasingly use AI to post vacancies, filter applications, score candidates, allocate tasks and monitor performance.

Zahed Ashkara guides organisations from AI Act obligation to review-ready evidence. This report captures that approach per domain, so your deployment is explainable and defensible towards supervisors and customers.

What this report is

- A practical readiness guide per domain or route
- An explanatory layer on the relevant AI Act articles
- A framework to build evidence

What this report is not

- Formal legal advice or a conformity assessment
- A replacement for DPIA, FRIA or vendor due diligence
- A guarantee of compliance

TABLE OF CONTENTS

What is inside this report.

01 Executive summary

02 Which HR applications fall under Annex III point 4

03 Provider or deployer: know which hat you wear

04 The Article 26 deployer obligations in practice

05 Bias, disparate impact and proving you do not discriminate

06 Transparency, explainability and the role of the works council

07 The review-ready HR-AI file

08 The obligations at a glance

09 Evidence pack: the checklist

10 Roadmap: 30, 60, 90 days

11 Frequently asked questions

12 Worksheets: supplier, classification and oversight

13 About the expert

KEY POINTS

AI in work and workforce management.

Employers, recruiters and HR-tech vendors increasingly use AI to post vacancies, filter applications, score candidates, allocate tasks and monitor performance. Under the EU AI Act (Regulation (EU) 2024/1689), many of these uses fall under Annex III point 4 and are in principle classified as high-risk AI. That triggers a heavy set of obligations for both the provider of the system and the employer that deploys it, and employers routinely underestimate that deployer role.

Status at 10 September 2026. Regulation (EU) 2026/1744 is in force. Chapter III, Sections 1 to 3 duties for high-risk AI under Article 6(2) and Annex III apply from 2 December 2027, except Article 6(5). Assess Article 111 transitional rules separately for existing systems. Article 4 and applicable prohibitions already apply; applicable Article 50 duties have applied since 2 August 2026. Paragraph 2 has a transition until 2 December 2026 only for systems placed on the market before 2 August 2026.

FINDING 01

Which HR applications fall under Annex III point 4

Annex III point 4 of the AI Act identifies two categories that count as high-risk.

FINDING 02

Provider or deployer: know which hat you wear

A provider develops an AI system or has it developed and places it on the market or puts it into service under its own name or trademark.

FINDING 03

The Article 26 deployer obligations in practice

Article 26 imposes a series of concrete obligations on the employer as deployer.

FINDING 04

Bias, disparate impact and proving you do not discriminate

Discrimination risk requires attention from both the provider and employer.

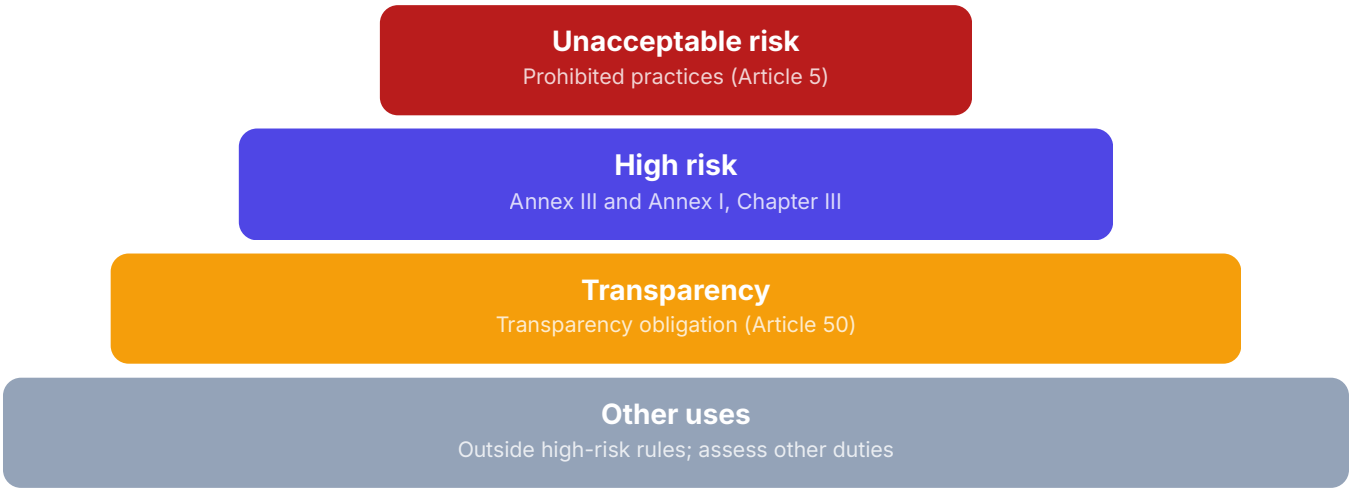
IN CONTEXT

Where you stand in the AI Act.

Annex III point 4	2 dec 2027	6	30 / 60 / 90
TOPIC	ANNEX III: CH. III, SECTIONS 1-3	POINTS TO ASSESS	PLANNING CADENCE IN DAYS

This concerns duties for systems under Article 6(2) and Annex III, except Article 6(5). First assess classification and any Article 111 transition. Articles 4, 5 and 50 may already apply.

THE AI ACT RISK PYRAMID



Schematic routes, not mutually exclusive categories: Article 50 can also apply to a high-risk system. Chapter V is a separate route for GPAI model providers. Start with actual use and your role.

THE REGULATION TIMELINE

2 feb 2025	Prohibited practices and AI literacy
2 aug 2025	GPAI model duties; see transition for existing models
27 jul 2026	Omnibus in force; amended AI literacy duty
2 aug 2026	Article 50 applies; GPAI enforcement active
2 dec 2026	Article 50(2) transition ends, only systems before 2 Aug 2026
2 aug 2027	GPAI models placed on the market before 2 Aug 2025
2 dec 2027	Chapter III, Sections 1-3: Article 6(2) / Annex III
2 aug 2028	Chapter III, Sections 1-3: Article 6(1) / Annex I

Chapter III: exception for Article 6(5). Also check Article 111 for each system. This overview is not a uniform deadline for all AI.

FROM OBLIGATION TO EVIDENCE

Classification › **Risk management** › **Documentation and logging** › **Human oversight** › **Evidence pack**

WHO CARRIES WHICH OBLIGATION

The following allocation concerns high-risk systems where the duty applies. Classification, role and transition determine the outcome.

Provider

Develops or has the system developed and places it on the market or puts it into service under its own name

- Chapter III: risk management, data and documentation (Art. 9 to 15)
- Conformity assessment and CE marking
- Registration under Art. 49: EU or national, where applicable

Deployer

Uses the system under its own authority

- Use per instructions and monitoring (Art. 26)
- Human oversight and log retention
- Fundamental rights assessment where required (Art. 27)

CHAPTER 02

Which HR applications fall under Annex III point 4

Annex III point 4 of the AI Act identifies two categories that count as high-risk. Sub a covers AI intended for the recruitment or selection of natural persons, in particular to place targeted job advertisements, to analyse and filter applications, and to evaluate candidates. Sub b covers AI intended to make decisions affecting the terms of work-related relationships, or relating to promotion, termination, task allocation based on individual behaviour or personal traits, or the monitoring and evaluation of the performance and behaviour of workers. In practice this means ATS filtering, CV scoring, video-interview analysis, matching algorithms, productivity and performance dashboards and automated task allocation almost all fall within scope.

Systems that filter, rank or evaluate candidates can be covered by point 4 when they support human decisions. The presence of a human does not automatically rule out high-risk status. Conversely, not every HR tool is high-risk: an administrative chatbot or writing aid is not listed in Annex III merely because of its HR context. Assess the specific intended use and then, where relevant, Article 6(3).

The legislator justifies this strict classification because HR-AI directly affects access to work and income, and therefore fundamental rights. Errors here are not abstract: a poorly trained filter can systematically disadvantage women, older candidates or people with a migration background without anyone noticing. The burden of proof therefore shifts to you: you must be able to demonstrate that the system does what it should, and does not do what it should not.

EXAMPLE FROM PRACTICE

A Dutch employer buys an international ATS with built-in CV ranking. HR assumes the vendor carries the responsibility. An internal audit reveals the employer qualifies as deployer under Article 26: there is no human oversight procedure, the works council was not informed and no bias test was run on the Dutch population. The tool was not the problem, the missing file was.

CHAPTER 03

Provider or deployer: know which hat you wear

A provider develops an AI system or has it developed and places it on the market or puts it into service under its own name or trademark. A deployer uses a system under its own authority in a professional context. Article 25 can transfer high-risk system provider status through rebranding, a substantial modification that leaves the system high-risk, or a change of intended purpose that makes a previously non-high-risk system high-risk. Routine configuration or fine-tuning does not automatically change the role; assess the actual modification.

For HR-tech vendors the position is reversed: they are typically providers and must ensure technical documentation, risk management, data quality, transparent instructions for use, logging by design, and a conformity assessment with CE marking and registration in the EU database. Employers who buy such a system rely on that documentation, but this does not relieve them of their own deployer obligations.

The first practical step is therefore a clear role and scope determination for each AI system in your HR stack. Who is the provider, who is the deployer, what is the intended purpose, and does it fall under Annex III point 4? Without that mapping you cannot build the rest of the file correctly, and you risk responsibilities falling through the cracks between you and your vendor.

CHAPTER 04

The Article 26 deployer obligations in practice

Article 26 imposes a series of concrete obligations on the employer as deployer. You must use the system in accordance with the provider's instructions for use (paragraph 1), assign human oversight to people with the necessary competence, training and authority (paragraph 2), and ensure input data is relevant and sufficiently representative for the intended purpose, to the extent you control it (paragraph 4). That last point is crucial in HR: if you feed the system historical recruitment data that is itself biased, you reproduce that bias.

For a risk within Article 79(1), Article 26(5) requires the deployer to inform the provider and market surveillance authority without undue delay and suspend use. For a serious incident, the deployer immediately informs first the provider, then the importer or distributor and the relevant authority. Record these routes alongside the provider’s own reporting duties and any sectoral notifications. Article 26(6) requires the deployer to retain automatically generated logs insofar as they are under its control. Choose an appropriate period of in principle at least six months, unless applicable Union or national law provides otherwise, including data protection law. Do not retain data longer than needed without a purpose or legal basis.

Two obligations matter specifically in the employment context. Paragraph 7 requires you, before deploying a high-risk AI system in the workplace, to inform worker representatives and affected workers that they will be subject to the use of the system, in line with applicable Union and national rules. And paragraph 11 requires you to inform affected persons that they are subject to a high-risk system, connected to the right to an explanation of individual decision-making under Article 86.

EXAMPLE FROM PRACTICE

A recruiter wants to use AI video analysis to evaluate candidates. First rule out prohibited inference of emotions from biometric data in a workplace context. For a permitted high-risk use, an authorised reviewer must be able to understand relevant output and, where appropriate, disregard or override it. Showing only a final score makes meaningful oversight difficult; an oversight procedure cannot legalise a prohibited use.

CHAPTER 05

Bias, disparate impact and proving you do not discriminate

Discrimination risk requires attention from both the provider and employer. Article 10 addresses high-risk provider data governance; Article 26 gives the deployer duties over input data under its control, monitoring and use. GDPR and equal-treatment law also continue to apply. Periodic bias analysis can support these duties, but the AI Act does not prescribe one fixed disparate-impact test for every employer.

Concretely, this means periodically measuring whether the outcomes of your filtering, scoring or matching differ systematically between groups, for example by gender, age or origin, to the extent that is lawful and proportionate to measure. You document the method used, the thresholds, the findings and the mitigating measures. A one-off test at go-live is insufficient: bias can creep in when data, population or model version changes.

The goal is a defensible file in which you can show that you proactively tested for disparate impact and intervened where needed. That moves your position from reactive, where you only explain what happened once a complaint arrives, to demonstrably in control. In a dispute or supervisory query, that is the difference between a short justification and a lengthy investigation with a possible fine via the supervisory authority.

CHAPTER 06

Transparency, explainability and the role of the works council

Organise information to workers and their representatives under Article 26(7). Article 86 provides, subject to its conditions, a clear explanation of the role of a high-risk Annex III system (except point 2) and the main elements of a decision based on it. The decision must produce legal effects or similarly significant adverse effects on health, safety or fundamental rights; statutory exceptions and other Union-law explanation rights remain relevant.

For Dutch practice, co-determination is added on top. Introducing or modifying an AI system that monitors, assesses or selects staff often triggers the consent or advisory rights of the works council under the Works Councils Act, particularly for arrangements around performance appraisal, recruitment and the processing of personal data. The AI Act information right and the works council route run in parallel and are best organised together.

In practice this means a transparency package: a clear explanation of what the system does and does not do, an explanation procedure for individual decisions, a communication line to candidates and workers, and a record of works council involvement. This is not only compliance, it is also trust: workers and candidates accept AI in HR considerably better when it is clear that a human decides and an explanation is available.

CHAPTER 07

The review-ready HR-AI file

All the individual obligations come together in a file that an auditor or supervisor can assess with little explanation. The backbone is an HR-AI register: an overview of every AI system used in recruitment, selection and workforce management, recording for each the provider, the deployer role, the intended purpose, the Annex III classification, the data flows used and the responsible owner within the organisation.

Around it you build the evidence layers: the instructions for use and how you follow them, the oversight procedure with designated and trained people, the bias and disparate-impact reports, the logging policy with the retention period of in principle at least six months, the DPIA and, where applicable, the FRIA under Article 27, the works council and worker communication, and the explanation procedure. Together this forms the narrative that you are in control.

A dossier must reflect working safeguards and system quality. Documentation helps reconstruct choices, tests, limitations and follow-up, but does not compensate for discriminatory outcomes or defective technology. Combine testing of actual behaviour with evidence of responsibility, human oversight and remediation.

The obligations at a glance

1

Annex III point 4 (a) and (b)

AI for recruitment and selection (targeted vacancies, filtering applications, evaluating candidates) and AI for decisions on working conditions, promotion, termination, task allocation based on behaviour or personal traits, and performance and behaviour monitoring is classified as high-risk AI.

2

Article 6(2) and 6(3)

Article 6(3) concerns only systems listed in Annex III. The assessment must establish no significant risk to health, safety or fundamental rights and at least one statutory condition: a narrow procedural task, improvement of completed human work, detection of decision patterns or deviations with proper human review, without replacing or improperly influencing the previous assessment, or preparation for a relevant assessment. Profiling of natural persons rules out the exception. The provider documents the assessment before market placement or putting into service under Article 6(4) and registers under Article 49(2). The deployer can request that reasoning from its supplier.

3

Article 26(2)

You assign human oversight to persons with the necessary competence, training and authority, with adequate support, so they can meaningfully assess and where necessary override AI-assisted HR decisions.

4

Article 26(6)

Article 26(6) requires the deployer to retain automatically generated logs insofar as they are under its control. Choose an appropriate period of in principle at least six months, unless applicable Union or national law provides otherwise, including data protection law. Do not retain data longer than needed without a purpose or legal basis.

5

Article 26(7)

Before deploying a high-risk AI system in the workplace, you inform worker representatives and the affected workers that they will be subject to the use of the system, in line with applicable Union and national rules and practice.

6

Article 27 and Article 86

Where applicable you carry out a fundamental rights impact assessment (FRIA), in particular public bodies and providers of public services, and you ensure that affected persons can obtain an explanation under Article 86 of the role of the AI system in a decision that significantly affects them.

Evidence pack: the checklist

- ✓ HR-AI register recording provider, deployer role, intended purpose and Annex III classification per system
- ✓ Role and scope determination (provider versus deployer) per AI system, including vendor contracts
- ✓ Oversight procedure with designated, trained owners and evidence of a genuine override capability
- ✓ Bias and disparate-impact reports with method, thresholds, findings and mitigation
- ✓ Logging policy and evidence of log retention for in principle at least six months
- ✓ Works council and worker communication (Article 26(7) plus co-determination) and explanation procedure (Article 86)
- ✓ DPIA under the GDPR and, where applicable, FRIA under Article 27
- ✓ Provider instructions for use and a record that actual use stays within them

Roadmap: 30, 60, 90 days

Indicative planning cadence. Feasibility depends on scope, existing measures and available evidence.

0 to 30 days

- Inventory all AI in recruitment, selection and workforce management and create a first HR-AI register
- Determine the role (provider or deployer) per system and test against Annex III point 4 and Article 6
- Identify the highest risks: automated pre-selection, scoring and performance monitoring
- Collect from vendors the instructions for use, conformity documentation and logging capabilities

30 to 60 days

- Set up the human oversight procedure with designated, trained owners (Article 26(2))
- Run a first bias and disparate-impact test and record the method and findings
- Start or update the DPIA and assess whether a FRIA (Article 27) is required
- Prepare the works council and worker communication and align the co-determination track (Article 26(7))

60 to 90 days

- Operationalise the logging policy with in principle at least six months retention and a monitoring routine
- Implement the explanation procedure for individual decisions (Article 86) and candidate and worker communication
- Bundle everything into a review-ready HR-AI file and run an internal gap review
- Assign ownership, a recurring bias-testing cycle and a procedure for new or modified HR-AI

Frequently asked questions

We only use purchased HR software. Are we still covered by the AI Act?

Buying HR software does not remove your own duties. First determine whether it is an AI system and which intended use is listed in Annex III point 4. For high-risk deployment you are usually the deployer with the applicable Article 26 duties, alongside the supplier’s duties. Not every procured HR feature is automatically high-risk.

Does this still apply if a human makes the final decision?

Usually yes. Annex III point 4 also covers systems that filter, rank or score candidates and thereby materially influence the human decision. A human in the loop does not automatically remove the system from the high-risk zone; instead it means you must demonstrate that the human control is meaningful.

What is the difference between a DPIA and a FRIA?

A DPIA under the GDPR focuses on risks to the processing of personal data. A FRIA under Article 27 of the AI Act focuses more broadly on the impact on fundamental rights and is in particular required for public bodies and providers of public services. They complement each other; the FRIA does not replace the DPIA.

Do we really have to inform the works council?

Yes. Article 26(7) requires you to inform worker representatives and affected workers before deploying a high-risk system in the workplace. In addition, Dutch co-determination law may trigger a consent or advisory right. Organise both tracks together.

When do these obligations take effect?

Status at 10 September 2026. Regulation (EU) 2026/1744 is in force. Chapter III, Sections 1 to 3 duties for high-risk AI under Article 6(2) and Annex III apply from 2 December 2027, except Article 6(5). Assess Article 111 transitional rules separately for existing systems. Article 4 and applicable prohibitions already apply; applicable Article 50 duties have applied since 2 August 2026.

What happens if we record nothing?

You stand weak in a complaint, dispute or supervisory query. Supervisors and courts assess demonstrability. Without a register, bias evidence, an oversight procedure and logging you cannot show you are in control, with a possible investigation and a fine via the supervisory authority as a result.

FROM INSIGHT TO ACTION

HR: supplier questionnaire

Use these questions before procurement or renewal. Request evidence for each answer. A sales claim or completed questionnaire is not a conformity assessment.

Field and prompt	Your answer / source + date
1. Purpose, scope and role Which HR decision does this system support? What version, intended use and limitations are documented? Who is the provider?	
2. Classification Which Annex III route have you assessed? If relying on Art. 6(3), what is your reasoning, and does the system profile people?	
3. Data and performance Which data, error rates and tests support use for our population? What evidence covers bias and limitations?	
4. Instructions and oversight How do users see uncertainty, errors and limitations? How can they disregard, correct or stop using outputs?	
5. Privacy and supply chain Which subprocessors, data flows, retention and training uses apply? What terms and evidence will you supply?	
6. Change and incident How do you report updates, changed purpose, risks and incidents? What documents/logs, audit rights and exit options do we receive?	

FROM INSIGHT TO ACTION

HR: classification sheet

Use one sheet per use case. Fill it in using a PDF reader or print it. Save your own copy and check that the answers were retained. For unknown facts, record who will obtain them and by when. Use file references, not personal data of affected people.

Field and prompt	Your answer / source + date
Use and decision Recruitment, selection, task allocation, evaluation, promotion or termination? Describe actual influence.	
AI definition and role Why is/is not this an AI system? Our role and supplier.	
Prohibited practice Check Art. 5 first, including workplace emotion recognition. Disclosure does not authorise prohibited use.	
Annex III point 4 Which exact use applies or does not apply, based on which facts?	
Possible Art. 6(3) route Obtain provider reasoning. For an Annex III system, profiling excludes the exception; record evidence and reviewer.	
Other duties and dates Assess GDPR/DPIA and any FRIA, staff information/ works council and transition. FRIA is not automatic for every HR use.	
Decision and follow-up Provisional conclusion, missing evidence, authorised reviewer, action owner and review date.	

FROM INSIGHT TO ACTION

HR: human oversight

Document the procedure before use. Where Art. 26(5) applies to a risk within Art. 79(1), prompt notification and suspension of use are mandatory.

Field and prompt	Your answer / source + date
Owner and decision-maker Who is accountable and who can override the output?	
Checks before use Which inputs, population and performance are checked?	
Checks for each output Which source and reasoning must the reviewer be able to see?	
Override and stop When must output be rejected or use suspended? Practical stop control and fallback process.	
Escalation Who receives a risk notice or incident? Which contact route and timing follow Art. 26 and your instructions?	
Records Where are review decisions, overrides, logs and instructions stored?	
Evaluation Which error/bias signals, review frequency and changes trigger reassessment?	

ABOUT THE EXPERT



Zahed Ashkara

AI Act consultant · AI governance specialist · lawyer and speaker

Zahed Ashkara advises organisations on the EU AI Act, GDPR and responsible AI adoption, bridging Legal, IT and Business. As a speaker and consultant he translates complex regulation into concrete, demonstrable steps: from classification to a review-ready evidence pack. He is the founder of Embed AI and Praxikon.

> **Compliance under the AI Act is not a technology question but a classification and evidence question. Start with the right classification, not the model.**

About Embed AI

Embed AI makes AI in organisations responsible and customer-ready under the EU AI Act and GDPR. From independent risk check and classification to a review-ready evidence pack, legal and technical, with implementation. For organisations with AI in their processes and vendors who sell AI.

Sources and next step

Version 2026.09 · Reference date 2026-09-10

This report references the official text of Regulation (EU) 2024/1689 and the European Commission guidelines. Read the official sources together with the amendments. Praxikon provides additional domain guidance.

EU AI Act: Regulation (EU) 2024/1689

Digital Omnibus on AI: Regulation (EU) 2026/1744, Article 1(39)-(40)

European Commission: implementation timeline

AI Act amendment: Regulation (EU) 2026/1744

Commission draft guidelines: high-risk classification (19 May 2026)

Praxikon · Current reports

Bring your HR use and the supplier questions that remain unanswered to the AI Act gap intake. Embed AI helps clarify the open point and determine an appropriate next step.

Discuss your open question