EU AI ACT · HIRING & HR

The AI in Hiring Compliance Starter Kit

A practical, plain-language starting point for HR and legal teams using AI to recruit.

By Signato · AI compliance, made clear

Current as of June 2026

Start here

You did not choose to become the AI Act person. Someone handed you a recruiting tool that screens CVs or ranks candidates, then handed you the question of whether it is legal, and the honest answer is that the rules landed faster than anyone wrote the playbook. The stakes are not small. Non-compliance with the high-risk obligations can reach 15,000,000 EUR or 3% of worldwide annual turnover, and the deadlines themselves keep moving, which means the ground shifts under you while you are still trying to stand on it.

This kit is the playbook for that person. It does not assume you have a room full of lawyers, because most mid-market teams do not. It turns the part of the EU AI Act that governs hiring into steps you can take this week and into evidence you can show when a regulator, a candidate's lawyer, or your own board asks to see your file. Every legal claim here carries its source, so you can check us and cite us.

One promise, plainly. The regulation is in motion, and a static PDF goes stale the moment a date changes. Signato exists to keep this current. We tell you what is settled, what is still open, and when the law actually moves we say so, instead of pretending we knew all along. Clear over clever, proof over promises, the standard you can actually meet.

How to use this kit

If your team uses software to advertise jobs, screen CVs, rank applicants, or score candidates, and that software uses artificial intelligence, this kit is for you. It is written for the person inside a mid-market company who has been handed this problem and does not have a room full of lawyers.

It does six things:

1. Helps you work out whether your use of AI in hiring is "high-risk" under the EU AI Act.
2. Gives you a checklist of the obligations that attach to that classification.
3. Provides a lean fundamental-rights and impact assessment template you can fill in.
4. Lists the questions to put to the vendor that sold you the tool.
5. Gives you an evidence log so you can show your work when someone asks.
6. Tells you, clearly, where this stops and a lawyer begins.

Each section ends with something you can actually do this week. Read the disclaimer at the end before you act on any of it.

One thing to know before you start. The EU AI Act timeline is moving right now. In late 2025 the European Commission proposed deferring the high-risk obligations, and in May 2026 the Parliament and Council reached a provisional political agreement to push the main high-risk date

from 2 August 2026 to 2 December 2027. As of the date on this edition, that change still needed formal adoption and publication in the Official Journal to take legal effect. We flag exactly where this matters below. When the regulation moves, we will say so rather than pretend we know. [Sources: European Parliament press release, 18 March 2026 (https://www.europarl.europa.eu/news/en/press-room/20260316IPR38219/meps-support-postponement-of-certain-rules-on-artificial-intelligence); Council of the EU press release, 7 May 2026 (https://www.consilium.europa.eu/en/press/press-releases/2026/05/07/artificial-intelligence-council-and-parliament-agree-to-simplify-and-streamline-rules/); Digital Omnibus on AI, European Commission (https://digital-strategy.ec.europa.eu/en/library/digital-omnibus-ai-regulation-proposal).]

Section 1. Is your AI hiring tool "high-risk"? A classification check

What the rule says

The EU AI Act (Regulation (EU) 2024/1689) sorts AI systems by risk. The category that matters most for hiring is **high-risk**. An AI system is high-risk when it falls within one of the use cases listed in **Annex III**, via **Article 6(2)**.

Annex III, **point 4 ("Employment, workers management and access to self-employment")** lists, verbatim:

> "(a) AI systems intended to be used for the recruitment or selection of natural persons, in particular to place targeted job advertisements, to analyse and filter job applications, and to evaluate candidates;
>
> (b) AI systems intended to be used to make decisions affecting terms of work-related relationships, the promotion or termination of work-related contractual relationships, to allocate tasks based on individual behaviour or personal traits or characteristics or to monitor and evaluate the performance and behaviour of persons in such relationships."

[Source: Annex III, Regulation (EU) 2024/1689 (https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng); Annex III text, AI Act Explorer (https://artificialintelligenceact.eu/annex/3/).]

In plain terms: if your tool helps decide **who gets advertised the job, who gets through the CV filter, who gets ranked or scored, who gets hired, promoted, or let go**, it is squarely inside the high-risk list.

Are you a "provider" or a "deployer"?

This determines which obligations are yours. The Act draws the line between the party that builds or sells the system and the party that uses it.

- **Provider:** the party that develops the AI system, or has it developed, and places it on the market or puts it into service under its own name or trademark (Article 3(3)).

- **Deployer:** the party that uses an AI system under its own authority, in the course of a professional activity (Article 3(4)).

[Source: Article 3 definitions, Regulation (EU) 2024/1689 (https://artificialintelligenceact.eu/article/3/).]

Most mid-market HR and legal teams are **deployers**: you bought a tool, you use it on your candidates. The vendor is usually the provider. This matters because deployer obligations are narrower than provider obligations, and because a chunk of your job is to make sure the provider has done theirs (that is what Section 4 is for).

One caution: under Article 25, a deployer can be treated as a provider in certain situations, for example if you put your own name on a high-risk system, substantially modify it, or change its intended purpose so that it becomes high-risk. If you are heavily customising a tool, do not assume you are only a deployer. [Source: Article 25, Regulation (EU) 2024/1689 (https://artificialintelligenceact.eu/article/25/).]

Is there a way out of "high-risk"?

Possibly, but treat it with care. Article 6(3) allows that a system in an Annex III area is **not** high-risk if it does not pose a significant risk of harm to health, safety, or fundamental rights, including where it only performs a narrow procedural task or improves the result of a previously completed human activity. **But** Article 6(3) expressly states that a system **is always high-risk** if it profiles natural persons. A provider that relies on this exemption must document its assessment before the system is placed on the market or put into service (Article 6(4)). [Source: Article 6, Regulation (EU) 2024/1689 (https://artificialintelligenceact.eu/article/6/).]

Be honest with yourself here. A tool that ranks or scores candidates is doing more than a narrow procedural task, and many hiring tools profile. Do not use Article 6(3) as a shortcut to wish the obligations away.

When does this bite?

This is the live question, and the honest answer has two layers: a date, and a warning not to trust the date too much. The original date for Annex III high-risk obligations to apply was **2 August 2026**. The Digital Omnibus provisional agreement of May 2026 would move that to **2 December 2027** for stand-alone Annex III systems (and 2 August 2028 for AI embedded in products regulated under Annex I). As of this edition, the change was provisionally agreed but not yet formally adopted and published, with adoption expected before 2 August 2026.

What this means for you: **do not plan around a single date as if it were settled.** Plan to be ready, confirm the operative date against the Official Journal before you rely on it, and watch for the formal adoption. A deferral is breathing room to do this properly, not a reason to do nothing. [Sources: Council of the EU, 7 May 2026 (https://www.consilium.europa.eu/en/press/press-releases/2026/05/07/artificial-intelligence-council-and-parliament-agree-to-simplify-and-streamline-rules/); Gibson Dunn analysis of the Omnibus agreement

(https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/); AI Act implementation timeline (https://artificialintelligenceact.eu/implementation-timeline/).]

> **Do this week:** Write one sentence for each AI tool your team uses in hiring, describing what decision it touches. Map each against Annex III point 4 (a) and (b) above. For each, note whether you are the provider or the deployer. That single page is the foundation for everything else in this kit.

Section 2. Checklist of the high-risk obligations (Annex III hiring)

These are the core requirements the EU AI Act attaches to high-risk systems. The first block sits primarily with the **provider** (Chapter III, Section 2, Articles 8 to 15). The second block is what you, as a **deployer**, owe directly (Article 26). Your practical task as a deployer is twofold: do your own obligations, and get evidence that the provider did theirs.

A. Provider obligations you should confirm have been met (Articles 8 to 15)

- ❑ **Risk management system** established and run across the system's lifecycle (Article 9).
- ❑ **Data and data governance:** training, validation and testing data sets are relevant, sufficiently representative, and examined for bias (Article 10).
- ❑ **Technical documentation** drawn up before the system is placed on the market and kept up to date (Article 11).
- ❑ **Record-keeping / logging:** the system automatically records events ("logs") over its lifetime (Article 12).
- ❑ **Transparency and information to deployers:** instructions for use that let you understand and use the system properly, including its capabilities and limitations (Article 13).
- ❑ **Human oversight** is designed into the system so a person can understand it, monitor it, and intervene (Article 14).
- ❑ **Accuracy, robustness and cybersecurity** at an appropriate level, declared and maintained (Article 15).

[Source: Chapter III Section 2, Articles 8 to 15, Regulation (EU) 2024/1689 (https://artificialintelligenceact.eu/chapter/3/).]

There is also a broader **quality management system** obligation on providers (Article 17), plus registration of the system in the EU database (Articles 49 and 71). You generally cannot verify these from the outside, so the practical move is to ask the vendor for evidence. See Section 4.
[Source: Articles 17, 49 and 71, Regulation (EU) 2024/1689 (https://artificialintelligenceact.eu/article/17/).]

B. Deployer obligations that are directly yours (Article 26)

- ❑ **Use the system in line with the provider's instructions for use** (Article 26(1)).
- ❑ **Assign human oversight** to people who have the competence, training, and authority to perform it (Article 26(2)).
- ❑ **Make sure input data is relevant and sufficiently representative** for the system's intended purpose, to the extent you control the input data (Article 26(4)).
- ❑ **Monitor the operation** of the system and, where you have reason to believe use creates a risk or a serious incident occurs, inform the provider and the relevant authority and suspend use as appropriate (Article 26(5)).
- ❑ **Keep the logs** automatically generated by the system, where they are under your control, for an appropriate period (Article 26(6)).
- ❑ **Inform workers and their representatives** before putting a high-risk system into use in the workplace, that they will be subject to it (Article 26(7)).
- ❑ **Inform individuals** when you use a high-risk system to make or assist decisions about them (linked to the affected-persons transparency duties; see Article 26(11) and Article 86 on the right to explanation).

[Source: Article 26, Regulation (EU) 2024/1689 (https://artificialintelligenceact.eu/article/26/); Article 86 (right to explanation of individual decision-making) (https://artificialintelligenceact.eu/article/86/).]

C. The fundamental rights impact assessment (Article 27)

A specific obligation worth calling out: certain deployers of Annex III high-risk systems must carry out a **fundamental rights impact assessment (FRIA)** before deployment. The duty in Article 27 applies to deployers that are bodies governed by public law, or private operators providing public services, and to deployers of certain banking and insurance use cases. Whether it strictly applies to a private-sector employer depends on the facts. [Source: Article 27, Regulation (EU) 2024/1689 (https://artificialintelligenceact.eu/article/27/).]

Our practical position: even where Article 27 may not strictly bind you, a lightweight impact assessment is the cleanest way to show you took fundamental rights seriously. Whether Article 27 is legally mandatory for your specific organisation is a question for your lawyer. Section 3 gives you a template to start the thinking either way.

> **Do this week:** Print the Article 26 list (block B). For each item, mark green (done and evidenced), amber (doing it but not written down), or red (not started). The amber items are your quickest wins: you are likely already doing the work and just need to record it.

Section 3. A lean fundamental-rights and impact assessment template

This is a starting structure, not the legal Article 27 form, and not a guarantee of compliance. It is designed to get a non-specialist team to think through the right questions and leave a record. Fill one in per high-risk hiring tool.

Note on Article 27. Where Article 27 applies to your organisation, the assessment must cover specific elements set out in that article (the deployer's process, the period and frequency of use, the categories of people affected, the specific risks of harm, human oversight measures, and the measures to take if risks materialise). Map your answers below to those elements, and have a lawyer confirm completeness. [Source: Article 27, Regulation (EU) 2024/1689 (https://artificialintelligenceact.eu/article/27/).]

AI Hiring Tool, Impact Assessment

1. The tool and the decision

- Tool name and vendor:
- What it does, in one sentence:
- Which hiring decision it touches (advertising / CV filtering / ranking / scoring / interview analysis / final selection):
- Provider or deployer (us):
- Annex III point 4 sub-point it maps to (a or b):

2. Who is affected

- Categories of people (external candidates, internal applicants, specific roles or regions):
- Estimated number of people per year:
- Any groups who could be disproportionately affected (consider sex, race or ethnicity, age, disability, and other protected characteristics):

3. How it is used

- Period and frequency of use:
- Does a human make the final decision, or does the tool decide? Describe the human's actual role, not the role on paper:
- What information does the candidate receive about the tool, and when:

4. Risks to fundamental rights

- Risk of discrimination or unequal treatment (and on what basis):
- Risk to privacy and data protection (note: this often triggers a separate GDPR Data Protection Impact Assessment; coordinate the two):
- Risk of error or false signals, and the consequence of a wrong output for a candidate:

- Risk of over-reliance by the human reviewer ("automation bias"):

5. Controls and oversight

- Who performs human oversight, and are they trained and empowered to overrule the tool:
- What the provider's bias testing showed, and when it was last done (cross-reference Section 4):
- Logging and record-keeping in place:
- Escalation path if something goes wrong:

6. Decision and sign-off

- Residual risk after controls (low / medium / high):
- Measures to take if a risk materialises:
- Decision to deploy / not deploy / deploy with conditions:
- Owner, date, and review date:
- Flagged for legal review (yes / no), and why:

> **Do this week:** Complete sections 1 and 2 of the template for your highest-volume hiring tool. Those two sections alone force the most important conversation: what the tool decides and who it affects.

Section 4. The right questions to ask your vendor (due diligence)

You depend on the provider having done its job. If they cannot answer these, that silence is itself a finding, and you should record it as one. Send this list, ask for written answers, and keep the answers in your evidence log.

On classification and role

1. Do you consider this system high-risk under Annex III of the EU AI Act? If not, on what basis (for example an Article 6(3) assessment), and can we see that assessment?
2. Do you act as the provider for this system? If we configure or brand it, does that change our role under Article 25?

On the core requirements (Articles 8 to 15) 3. Can you share the **instructions for use** required under Article 13, including the system's intended purpose, capabilities, and known limitations? 4. What **bias and representativeness testing** have you done on the training, validation, and testing data (Article 10), and what were the results? When was it last run? 5. How does the system support **human oversight** (Article 14) in practice, and what can a reviewer see and override? 6. What are the system's stated **accuracy, robustness, and cybersecurity** levels (Article 15), and how are they measured? 7. Do you maintain **technical documentation** (Article 11) and automatic **logs** (Article 12), and what can you provide to us?

On conformity and registration 8. Has the system undergone the required **conformity assessment**, and is there an **EU declaration of conformity** and **CE marking**? Can we see them? (Articles 43, 47, 48.) 9. Is the system **registered in the EU database** for high-risk systems (Articles 49 and 71)?

[Sources: Articles 10 to 15, Regulation (EU) 2024/1689 (https://artificialintelligenceact.eu/chapter/3/); Articles 25, 43, 47, 48, 49, 71 (https://artificialintelligenceact.eu/).]

On overlapping local laws (ask if you hire in these places) 10. **New York City:** If we use this tool for roles in NYC, has it had an independent **bias audit** within the last year, and will you provide the published results so we can meet Local Law 144? (See Section 5 note.) 11. **United States generally:** What do you provide to help deployers meet emerging US state AI and automated-decision laws (for example transparency and notice obligations)?

> **Do this week:** Send questions 3, 4, and 8 to the vendor of your most-used tool. The instructions for use, the bias testing results, and the declaration of conformity are the three documents that do the most to show your file is in order.

Section 5. Evidence log template

If you cannot show it, you cannot rely on it. The point of this log is that, on the day someone asks ("a regulator, a candidate's lawyer, your own board"), you can produce a dated record instead of a memory. Keep it simple and keep it current.

AI Hiring Compliance, Evidence Log

#	Date	Tool / vendor	Obligation or question (article)	Evidence held (document, link, owner)	Status (green / amber / red)	Next action and due date
1			Annex III classification note (Art. 6, Annex III)			
2			Provider vs deployer determination (Art. 3, 25)			
3			Instructions for use received (Art. 13)			
4			Bias / data testing results (Art. 10)			
5			Human oversight assignment and training (Art. 14, 26(2))			
6			Logs retained (Art. 12, 26(6))			
7			Worker / representative notice (Art. 26(7))			

#	Date	Tool / vendor	Obligation or question (article)	Evidence held (document, link, owner)	Status (green / amber / red)	Next action and due date
8			Candidate notice and explanation (Art. 26, 86)			
9			Impact assessment completed (Art. 27 / internal)			
10			Declaration of conformity / CE / EU database (Art. 47, 48, 49)			
11			Vendor due-diligence answers on file (Section 4)			

[Article references drawn from Regulation (EU) 2024/1689 (https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng).]

A note on penalties, so the stakes are clear, not to frighten you. The EU AI Act sets administrative fines in tiers under Article 99. Breaching the prohibited-practices rules (Article 5) can reach up to 35,000,000 EUR or 7% of total worldwide annual turnover, whichever is higher. Non-compliance with other obligations, including the high-risk obligations on providers and deployers, can reach up to 15,000,000 EUR or 3% of worldwide annual turnover, whichever is higher. Supplying incorrect, incomplete, or misleading information to authorities can reach up to 7,500,000 EUR or 1%, whichever is higher. For SMEs and start-ups, Article 99 applies the **lower** of the amount or the percentage. The actual fine in any case depends on the authority's assessment of factors set out in the Act. [Source: Article 99, Regulation (EU) 2024/1689 (https://artificialintelligenceact.eu/article/99/); EUR-Lex full text (https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng).]

A note on local US laws, because they can apply today even where the EU date has moved.

- **New York City, Local Law 144 of 2021.** If you use an "automated employment decision tool" for candidates or employees in NYC, the law requires an independent **bias audit** within one year of use, **publication** of a summary of the audit results, and **notice** to candidates and employees. It took effect on 1 January 2023, with enforcement by the Department of Consumer and Worker Protection (DCWP) from 5 July 2023. DCWP can impose civil penalties of 500 to 1,500 USD per violation, and ongoing violations can accrue daily. [Sources: DCWP, Automated Employment Decision Tools (https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page); NYC Administrative Code, Subchapter 25 (https://codelibrary.amlegal.com/codes/newyorkcity/latest/NYCadmin/0-0-0-135598).]
- **Colorado AI Act (SB 24-205).** Originally a comprehensive, risk-based law for high-risk AI used in consequential decisions including employment. Its effective date has been delayed more than once, and in May 2026 the Governor signed **SB 189**, which delayed the effective date to **1 January 2027** and significantly narrowed the law: it removed the deployer duty of care, the risk-management program, and the impact-assessment requirements, moving to a narrower focus on disclosure and transparency for certain automated decision-making technologies. **Do not**

build a Colorado program off the original 2024 text; confirm the current SB 189 version before relying on it. [Sources: Hunton, Colorado AI Act Amended and Effective Date Delayed (https://www.hunton.com/privacy-and-cybersecurity-law-blog/colorado-ai-act-amended-and-effective-date-delayed); Colorado General Assembly, SB24-205 (https://leg.colorado.gov/bills/sb24-205).]

> **Do this week:** Create the log as a shared spreadsheet and fill in rows 1, 3, and 4 for your main tool, even if the status is red. An honest red with a due date is worth more than a blank.

Where this goes next

You are now on the Signato list, and that is deliberate, because a compliance document with a fixed date on it starts decaying immediately. The free weekly intelligence does the opposite: it tracks the AI Act and the laws around it, tells you when something actually moves (not when a headline merely speculates), and is segmented by role, so the HR reader and the legal reader each get what is theirs and not a wall of the other's detail. The deferral covered in Section 1 is exactly the kind of shift you will hear about from us first, with the operative date and the source, the moment it is real.

The kit in your hands is the orientation layer. The work itself lives in documents you fill in, revise, and defend, and that is what the Signato **Professional** tier is being built to give you: editable, working versions of the templates referenced throughout this kit. The full fundamental rights impact assessment, the model card, the risk register, and the vendor due-diligence questionnaire, as files you can complete and adapt, not static pages you retype. Every template carries a version stamp, so when an article or a date changes you can see at a glance what is current and what needs a second look, and role-based alerts tell the right owner on your team when their part has moved.

Professional is arriving soon, and there is no checkout to rush you toward today. If you want it the moment it opens, reply to your Signato welcome email with the single word **Pro**. That puts you on the priority waitlist, ahead of the general release, with no commitment attached. We would rather earn the upgrade than sell you a link.

Section 6. Disclaimer

This Starter Kit is an educational starting point, not legal advice. Signato is not a law firm and nothing here creates a lawyer-client relationship. The EU AI Act and the local laws referenced here are complex, fact-specific, and changing, and several key dates were unsettled when this was written. Nothing in this kit guarantees compliance, and following it will not on its own make any tool, vendor, or organisation compliant.

Use this kit to get oriented, to organise your evidence, and to have a sharper conversation with a qualified lawyer in the relevant jurisdiction before you make decisions or rely on any classification, deadline, or obligation described here. Where this kit points to a specific risk or an unsettled question, that is exactly where professional advice is worth getting.

Every regulatory statement in this document carries its source. If you find a point where the law has moved since the date on this edition, treat the source, not this document, as the authority.
