zortify

White paper

Responsibilities & Best Practices of HRs under the EU AI Act v2.0

For HR leaders who buy and deploy third-party AI tools

1. Executive Summary — Practical, human, and audit-ready

AI now sits at the heart of modern HR — from screening to development. The EU's approach has matured too: from the **2020 AI White Paper** to the **EU AI Act** (Reg. 2024/1689), which introduces a phased, risk-based framework through 2026—2027. A voluntary **Code of Practice for General-Purpose AI (GPAI)** arrived on **10 July 2025**, guiding transparency, copyright, and safety for model providers.

This paper focuses on a very specific (and common) scenario: **HR as the deployer of third-party AI services**. You don't build the models — but you **are** accountable for how they're used. The good news: when technology control is limited, **organizational controls** (clear processes, vendor diligence, human oversight) do the heavy lifting. Apply the guidelines here and you'll be able to demonstrate compliance — calmly — in internal and external audits.

What you'll get: a simple vendor checklist, governance by company size, a no-drama incident plan, and how Zortify supports you end-to-end.

2. Regulatory Context — What changes when, and why it matters

- **In force:** The AI Act was published on **12 July 2024** and **entered into force on 1 August 2024**.
- **GPAI Code of Practice:** Published **10 July 2025**; it's **voluntary** but designed to help model providers meet obligations.
- **Application timeline:** Core rules start applying broadly by **2 August 2026**, with staggered provisions before and after; the framework is effectively **fully operational by 2027**.

Why HR should care now: High-risk HR use cases (e.g., recruiting, evaluation) carry strict duties around **data quality, transparency, human oversight, and security**. And while vendors must do a lot, **deployers (you)** must prove that the AI is used responsibly — in context — inside your organization. Zortify calls this out plainly: **HR must deliver on AI — but not alone.**

3. Shared Responsibilities — Who owns what (in plain language)

3.1 HR as the deployer (that's you)

- **Classify correctly:** Tools used in hiring and performance management typically fall under **high-risk**. Treat them accordingly.
- **Be transparent:** Notify candidates/employees, and align **privacy policies** with how the tool actually processes data.
- **Keep humans in charge:** Document **human-in-the-loop** review and the ability to **override** automated recommendations.
- **Log for traceability:** Retain machine + human decision logs long enough to answer questions later.
- **Raise AI literacy:** Train HR on the tool's capabilities, limits, and oversight roles. (HR doesn't need to code; HR needs to ask the right questions.)

3.2 What to ask providers (and why)

Ask for a concise package you can file and use:

- **EU AI Act compliance documentation** incl. **TOMs** (bias mitigation, robustness, security, oversight).
- **Technical documentation** on data governance, bias testing, limitations.
- **Transparency assets** (plain-language explainer text you can share with candidates/employees).
- **Human oversight features** (review queues, rationale views, override paths).
- **Privacy alignment** (data types, retention, deletion, subprocessors, cross-border transfers).
- **Voluntary Code of Practice (GPAI) statement**, if relevant for the underlying model.

At a glance

> If a provider anticipates your compliance needs — with documentation, training, and support — you've found a partner, not just a tool.

4. The HR Quick-Check (use this with every AI vendor)

Ask the provider for…	Then verify inside HR…
Compliance docs incl. TOMs	Do they clearly explain **bias, security, oversight** and update regularly?
Technical doc + limitations	Are **known limits** clear so you can plan human review where it matters most?
Transparency notices	Are messages **plain-language** and consistent with your HR communications?
Human oversight workflow	Do you have a **named reviewer** and a **documented override** process?
Privacy alignment	Is your **privacy policy** updated and accurate?
GPAI Code of Practice stance	If applicable, is the provider proactively aligning?

5. Governance & Best Practices — Right-sized, not red-taped

5.1 Scale the setup to your company size

Micro enterprises (<10) — often startups

- **Committee:** HR lead + CTO/founder (legal counsel as needed)
- **Process:** Light AI inventory, vendor declarations, short trainings, annual review
- **Artifacts:** One tracker, self-declarations, meeting notes

Small (<250)

- **Committee:** HR Director, Legal/Compliance, IT, employee rep
- **Process:** Formal inventory, annual fairness reviews, vendor due diligence (biannual preferred)
- **Artifacts:** Policies, certificates, audit-ready minutes

Medium (≈3000) & Large (3000+)

- **Committee:** HR Tech, Legal, IT Security, DPO, CTO
- **Process:** Biannual cycle, third-party audits, advanced monitoring; 1—3 FTE for AI governance

5.2 Must-do activities (all sizes)

- Maintain an **AI tool inventory** with use case, risk class, and vendor evidence.
- Run **fairness/accuracy checks** at least annually (twice per year preferred), especially for recruiting and performance.
- Disclose AI use via **privacy policies** and candidate/employee notices.
- Enforce **human oversight** with clear roles and escalation paths.
- Keep **decision logs** and **training records** — audit-ready.

6. Common Pitfalls — and how to avoid them

- **"Vendor compliant = we're done."** Not quite. Deployers are accountable for **how** AI is used in context.
- **Opaque communications.** If people don't know AI is used, trust erodes.
- **No human override.** High-risk HR use needs meaningful human review.
- **Policy drift.** New tool, same policy? Update it.
- **Confusing GPAI guidance with high-risk rules.** They're related but **not the same**. GPAI Code helps model providers; your HR use is still high-risk.

7. Incident Response & Remediation — Calm beats chaos

Incidents happen. Your plan prevents escalation and speeds recovery.

7.1 Typical HR-AI incidents

- **Data breach/leakage:** Personal data exposure through interface or model behavior.
- **Discriminatory outcomes:** Systematic, **unjustified** disparities across protected groups.

- **System malfunction:** Crashes, nonsense outputs, corrupted rankings.
- **Privacy violations:** Processing beyond scope, missed deletion requests, unlawful transfers.

7.2 First 2 hours → 24 hours → 7 days

- **0—2 hours: Pause** the tool if needed, **preserve logs**, notify Legal/Compliance/HR leadership.
- **Within 24 hours:** Assess scope and personal-data impact; if relevant, prepare **GDPR notifications** (72-hour clock). Contact the provider via **pre-agreed incident channel**.
- **1—7 days:** Joint investigation with the provider, root cause + fix, review affected decisions, document everything.

7.3 Build it into the contract

- **Mandatory notification:** Discovery → notify within **2—4 hours**; full report in **24—48 hours**.
- **Shared investigation:** Access to relevant logs/data; named technical owner; **RTO/RPO** aligned to criticality.
- **Cost allocation:** If provider fault, they cover re-screening/legal costs tied to their negligence.
- **24/7 reachability:** A practical emergency contact (scaled to provider size).

8. How Zortify Implements This (so you don't have to start from zero)

- **Right-sized high-risk management:** Clear, **human-centered** workflows that keep people — not machines — in charge.
- **Complete documentation pack:** TOMs, transparency notices, privacy-policy language — ready for audits.
- **Explainability & oversight:** Review paths, rationale views, and override capabilities built for HR teams.
- **Early alignment with GPAI Code of Practice:** We support the Code's spirit (transparency, safety) even though GPAI obligations target model providers.
- **Education for HR:** Practical trainings and workshops; **AI literacy without the math**

class.

- **Incident readiness:** ISO-aligned RPO/RTO targets; named Customer Success contacts; clear escalation.

9. FAQ — Fast answers for busy HR leaders

Q1. Does this apply to in-house AI? No — this paper covers HR teams **buying and deploying** third-party tools.

Q2. What belongs in the privacy policy? Disclose AI use, data categories, purposes, rights (incl. explanation and human review), retention, and contacts.

Q3. What are TOMs? Technical & Organizational Measures that cover **data quality/bias, security, incident management, and oversight roles**.

Q4. What is the July 2025 GPAI Code of Practice? A **voluntary** framework for model providers to align with AI Act goals on transparency, copyright, and safety.

Q5. How often should we review outcomes? At least **annually** (ideally twice a year) for fairness, accuracy, and user feedback.

Q6. Can we deploy tools before "full" AI Act applies? Use compliant providers and **document** safeguards and progress. High-risk HR tools must meet requirements as relevant provisions come into force (broad application from **August 2026**; full operability by **2027**).

Q7. What if a provider shutters a product? Use **data portability** clauses, keep local copies of key documents, and maintain an **exit plan**.

Q8. Do we need separate consent? Depends on context (legitimate interests/contract may apply). Seek legal advice for your jurisdiction and use case.

Q9. Multi-country HR data? AI Act applies EU-wide; ensure lawful transfers and name a lead supervisory authority where needed.

Q10. High-risk vs. GPAI — what's the difference? High-risk is about your **use case** (e.g., recruiting). **GPAI** rules are for **model providers**. Both matter, differently.

Q11. Can employees opt out? They must have **meaningful human review** and the right to contest AI-influenced decisions.

Q12. How long to keep logs? Keep decision and oversight logs long enough to cover audits and legal limitation periods (often **3—7 years**; confirm locally).

10. References

- **EU AI Act (Reg. 2024/1689):** Official Journal & entry-into-force
- **AI Act implementation timeline (EPRS/EP analysis):** application from **2 Aug 2026**, fully effective by **2027**
- **GPAI Code of Practice (10 July 2025):** European Commission resources
- **Zortify perspective for HR leaders:** EU AI Act + HR competence and provider checklists

zortify

Zortify S.A.
9, rue du Laboratoire
L-1911 Luxembourg
https://zortify.com
hello@zortify.com