

# Global AI Risk & Regulation

Researched by Hey Lefty

Automated research briefings on topics you choose — heylefty.com

## TL;DR

Global AI governance has undergone a major structural realignment as the European Union officially deferred its high-risk compliance deadlines while tightening bias-screening and registration rules May Summary. Meanwhile, the United Kingdom has abandoned its voluntary stance by enacting its first statutory mandate for an AI and automated decision-making code of practice UK Enacts SI 2026/425. These shifts, paired with rapid US consumer protection litigation, signal that while compliance timelines are expanding, the legal boundaries for automated systems are rapidly hardening.

## The European High-Risk Deferral and Substantive Tightening

European regulators are trading immediate enforcement deadlines for stricter substantive compliance standards.

*"...the final agreement reinstates a strict necessity standard... processing special-category personal data for bias detection and correction..." — Gibson Dunn*

*"Providers seeking exemption from high-risk classification... will still need to register them in the EU database for high-risk systems, albeit with reduced information requirements. This represents a reversal from earlier drafts..." — Travers Smith*

While enterprises receive temporary operational breathing room from deferred deadlines pushing high-risk obligations to 2027 and 2028, the return to a strict necessity standard for bias screening means compliance teams cannot easily process sensitive personal data without exhaustive justification EU AI Act Omnibus Agreement. This shifts the enterprise focus from rushing deployment to executing rigorous, long-term technical audits. Furthermore, risk teams must note that although deadlines for high-risk systems are pushed back, basic transparency and disclosure rules are still set to go live in 2026, meaning compliance work cannot be paused May Summary.

**What to watch:** Whether organizations use this deferred timeline to build the necessary technical infrastructure for the looming transparency obligations, which remain active starting in 2026 May Summary.

## The UK's Hard Pivot to Statutorily Backed Data Governance

The United Kingdom is abandoning its voluntary, sector-led approach to artificial intelligence in favor of a mandatory, data-driven statutory regime.

*"The Commissioner must prepare an appropriate code of practice giving guidance as to good practice in the processing of personal data under the relevant data protection legislation in relation to— (a) developing and using artificial intelligence, and (b) automated decision-making." — SI 2026/425*

*"The code of practice must include guidance as to good practice in the processing of children's personal data." — SI 2026/425*

By legally forcing the Information Commissioner's Office to establish a statutory Code of Practice in 2026, the UK is transforming persuasive guidance into a powerful enforcement tool that courts will use to penalize non-compliant profiling UK Enacts SI 2026/425. This creates an immediate risk for companies using automated systems within the UK, particularly those processing children's data Fieldfisher. Consequently, any enterprise relying on automated decision-making for marketing, personalization, or user profiling must prepare for a much stricter regulatory environment May Summary.

**What to watch:** How the Information Commissioner's Office structures its upcoming public consultation on the draft code, which will signal the exact compliance boundaries for automated profiling May Summary.

## The Dual-Front Compliance Squeeze on Automated Recruiting

Automated hiring platforms face a dual-front compliance squeeze as European regulatory delays clash with immediate, aggressive litigation under consumer protection laws in the United States.

*"The Eightfold case isn't another AI discrimination lawsuit. It's a consumer protection action that reframes how plaintiffs can attack automated hiring." — Jones Walker LLP via Eightfold AI Class Action*

*"...the court may need to resolve that the AI vendor may not qualify as 'consumer reporting agencies' because it arguably does not assemble or evaluate information 'for the purpose of providing consumer reports to third parties,' as required by the statute." — Epstein Becker Green via May 28 Update*

While European recruiters get a temporary reprieve on high-risk obligations, they cannot escape accountability, as even exempted procedural tools require formal registration in the EU database EU AI Act Omnibus Agreement. Meanwhile, US courts are moving rapidly, meaning enterprise legal teams cannot treat automated hiring as a future compliance problem; they must address immediate contractual and statutory exposures today May 28 Update. The strategy of reframing candidate scoring as unauthorized credit reporting threatens to bypass traditional vendor contract protections, leaving employers exposed to massive statutory damages while European compliance structures are still being built Eightfold AI Class Action.

**What to watch:** Whether US courts allow the consumer protection framing to bypass traditional algorithmic bias defense strategies in upcoming dismissal hearings Eightfold AI Class Action.

## What surprised us

- **The Bias-Testing Dilemma Reinstated.** In a shocking reversal of earlier drafts, the EU Omnibus restored a "strict necessity" test for processing special category personal data like race or gender to detect and correct algorithmic bias EU AI Act Omnibus Agreement. This means developers face a legal dilemma: they are legally mandated to prevent bias, but processing the very data needed to test for it is heavily restricted unless they can prove no "less intrusive" alternative exists Travers Smith.
- **No Quiet Escapes for "Exempt" Systems.** Providers attempting to carve out their AI from high-risk classification—by claiming their tools only perform narrow procedural tasks—will still be forced to register them in the public EU database EU AI Act Omnibus Agreement. The EU has effectively closed the door on stealth deployments of borderline high-risk systems Travers Smith.
- **The Death of the UK's "Light Touch" Era.** The enactment of new regulations officially ends the UK's long-standing, voluntary "sector-led" regulatory approach UK Enacts SI 2026/425. By forcing the ICO to write a statutory Code of Practice, the UK is quietly building a hard regulatory framework backed by the full weight of GDPR enforcement May Summary.

### Open threads worth a vote

- **Eightfold AI Motion to Dismiss: FCRA/ICRAA Precedent for Automated Hiring Tools** — Cast your vote to track whether the federal court's upcoming hearing will rule that automated candidate scoring constitutes "consumer reporting" under the Fair Credit Reporting Act, establishing a critical liability precedent for algorithmic recruiting tools May 28 Update.

## Appendix: Findings

# May 29, 2026 Cycle Summary: European AI Regulatory Realignment — The EU's High-Risk Deferral and the UK's First Statutory AI Code

This research cycle surfaced two monumental developments in European AI liability, compliance, and governance that collectively signal a major structural realignment in how both the EU and the UK are regulating artificial intelligence. While the EU is extending compliance timelines for high-risk systems to match the reality of its lagging regulatory infrastructure, the UK is taking its first major step toward statutorily backed, data-driven AI rules.

Together, these changes create a dual-track landscape for enterprise legal and risk teams: a temporary "breathing room" on EU high-risk obligations, contrasted with a tightening of data protection enforcement and profiling rules in the UK.

### 1. The EU AI Act Omnibus: Pragmatic Deferral and Substantive Tightening

Following the provisional political agreement on the **Digital Omnibus on AI** on May 7, 2026 (confirmed by Member States on May 13), the EU has deferred the applicability of high-risk AI system (HRAIS) obligations:

- **Annex III Stand-alone HRAIS** (e.g., recruitment, credit scoring) is postponed to **December 2, 2027**.
- **Annex I Embedded HRAIS** (e.g., medical devices, machinery) is postponed to **August 2, 2028**.

However, this timeline relief comes with a tightening of substantive rules:

- **Strict Bias screening:** Reversing earlier proposals for flexible bias-testing rules, the final agreement reinstates a strict necessity test for processing special category data.
- **Exempted Systems Registration:** Providers seeking to bypass high-risk classification under Article 6(3) must still register in the EU database.
- **New Prohibitions:** A new ban on non-consensual sexually explicit content ("nudifiers") is introduced under Article 5.
- **Reinforced AI Office:** The EU AI Office is granted exclusive competence over GPAI-integrated systems and equipped with aggressive new enforcement tools, including on-site inspections and direct fines.

*For more details, see:* [EU AI Act Omnibus Agreement: Definitive High-Risk Deferral, Strict Bias Screening, and Expanded AI Office Powers Enacted]

### 2. The UK's First Statutory AI Code: Enactment of SI 2026/425

In a major pivot from its historically voluntary, sector-led AI policy, the UK has enacted **The Data Protection Act 2018 (Code of Practice on Artificial Intelligence and Automated Decision-Making) Regulations 2026 (SI 2026/425)**, which came into force on **May 12, 2026**.

This statutory instrument places a binding legal duty on the Information Commissioner's Office (ICO) to draft and publish a statutory Code of Practice for processing personal data in the context of AI and automated decision-making (ADM). The code must include specific, rigorous guidelines on processing children's data. Once finalized and approved by Parliament, this code will carry statutory weight, transforming the ICO's previously voluntary guidance into a powerful enforcement tool that courts and regulators will use to assess UK GDPR compliance.

*For more details, see:* [UK Enacts SI 2026/425: Mandating the First Statutory Data Protection Code of Practice for AI and Automated Decision-Making]

### **Enterprise Risk and Compliance Takeaways**

- **Do Not Pause EU Compliance:** August 2, 2026, remains a live compliance date for Article 50 transparency obligations. Use the HRAIS deferrals to build robust compliance frameworks rather than pausing efforts.
- **Prepare for UK ADM Scrutiny:** Any enterprise utilizing automated profiling, candidate screening, or algorithmic decision-making under UK jurisdiction must prepare for the ICO's upcoming draft code, which will heavily scrutinize automated decisions and children's data.
- **Evaluate Small Mid-Cap (SMC) Status:** EU-operating enterprises should assess whether they qualify for the newly extended SMC relaxations to reduce their administrative compliance burden under the AI Act.

## **UK Enacts SI 2026/425: Mandating the First Statutory Data Protection Code of Practice for AI and Automated Decision-Making**

In a major departure from its historically voluntary, sector-led approach to artificial intelligence regulation, the United Kingdom has established its first statutory obligation for an AI and Automated Decision-Making (ADM) code of practice. On **May 12, 2026**, **The Data Protection Act 2018 (Code of Practice on Artificial Intelligence and Automated Decision-Making) Regulations 2026 (SI 2026/425)** officially entered into force across England, Wales, Scotland, and Northern Ireland.

This statutory instrument places a binding legal duty on the UK Information Commissioner's Office (ICO) to draft, consult on, and publish a statutory Code of Practice that will govern how personal data is processed when developing and deploying AI and automated decision-making systems.

### **1. Key Provisions of SI 2026/425**

- **Mandatory Code of Practice:** Under Regulation 2(1), the Information Commissioner *must* prepare a code of practice providing guidance on "good practice" in the processing of personal data under the UK GDPR and the Data Protection Act 2018 (excluding intelligence services processing under Part 4) in relation to:
  1. Developing and using artificial intelligence, and
  2. Automated decision-making (ADM).
- **Focus on Children's Data:** Regulation 2(2) explicitly mandates that the statutory code must include "guidance as to good practice in the processing of children's personal data."
- **Defining Automated Decision-Making:** ADM is defined with specific reference to Article 22C(1) of the UK GDPR and section 50C(1) of the 2018 Act, connecting the code directly to statutory restrictions on profiling and automated individual decision-making.
- **National Security Exemption:** Regulation 3 modifies Section 124B of the Data Protection Act 2018, specifying that the advisory panel established to review the Commissioner's draft code "must not consider or report on any aspect of the code relating to national security."

### Verbatim Quotes

From the statutory text of SI 2026/425:

"The Commissioner must prepare an appropriate code of practice giving guidance as to good practice in the processing of personal data under the relevant data protection legislation in relation to— (a) developing and using artificial intelligence, and (b) automated decision-making."  
 "The code of practice must include guidance as to good practice in the processing of children's personal data."

From the Fieldfisher Legal Update:

"The ICO is the next focus and the team looks at... a new Code of Practice on AI and Automated Decision Making Regulations together with the new guidance on storage and access technologies."

### What This Means for Enterprises

While SI 2026/425 does not directly regulate AI software or impose direct civil liability on its own, it creates the legal mandate for a statutory "rulebook" that the ICO will use to enforce the UK GDPR.

- **Transition from Guidance to Statutorily Backed Rules:** Historically, the ICO's AI guidance and risk toolkits were persuasive but voluntary. Once the statutory Code of Practice is finalized and approved by Parliament, it will carry significant weight in courts and regulatory enforcement actions. Non-compliance with a statutory code under the Data Protection Act 2018 is highly influential in establishing whether an organization breached the UK GDPR.
- **Heightened Focus on AI Profiling and Children:** Any enterprise using AI for automated hiring, credit scoring, marketing profiling, or user personalization—particularly where children are involved—must prepare for a much stricter, statutorily backed regulatory regime.
- **Upcoming Public Consultation:** The ICO is expected to launch a broad public consultation on the draft code in mid-2026, which will serve as a key opportunity for enterprises to provide input on the operational feasibility of the proposed rules.

#### Sources

- The Data Protection Act 2018 (Code of Practice on Artificial Intelligence and Automated Decision-Making) Regulations 2026 (SI 2026/425)
- Data and Privacy Matters: Legal Updates - April 2026

## EU AI Act Omnibus Agreement: Definitive High-Risk Deferral, Strict Bias Screening, and Expanded AI Office Powers Enacted

In a major regulatory recalibration, the European Union institutions reached a provisional political agreement on the "Digital Omnibus on AI" on May 7, 2026, which was subsequently confirmed by Member State representatives in the Council on May 13, 2026. This package provides critical compliance relief for enterprises by postponing the applicability dates for high-risk AI systems (HRAIS) by over a year, while simultaneously introducing strict new prohibitions, rolling back earlier proposals that would have relaxed bias-screening rules, and expanding the enforcement powers of the central EU AI Office.

While the core architecture of the EU AI Act remains intact, the Omnibus represents a pragmatic adjustment to the fact that the technical standards and regulatory infrastructure required to make the Act operable have not materialized on the original schedule.

### 1. The Postponed HRAIS Compliance Timeline

The Omnibus replaces the European Commission's originally proposed conditional trigger mechanism with fixed, deferred deadlines:

- **Annex III Stand-alone HRAIS (Recruitment, Credit Scoring, Education, Law Enforcement):** Applicability is postponed by 16 months to **December 2, 2027** (originally August 2, 2026).
- **Annex I Embedded HRAIS (Medical Devices, Toys, Machinery, Vehicles):** Applicability is postponed by 12 months to **August 2, 2028** (originally August 2, 2027).

Importantly, August 2, 2026, remains a live compliance date for Article 50 transparency obligations (e.g., disclosing AI interaction), though existing systems receive a four-month grace period (until December 2, 2026) to implement watermarking under Article 50(2).

## 2. Key Substantive Overhauls and Compromises

- **Strict Bias Screening Standards:** In a major reversal of earlier drafts that proposed more flexible rules for processing special category personal data for bias detection, the final agreement reinstates a strict necessity standard.
- **Registration of Exempted High-Risk Systems:** Providers seeking to carve out their systems from high-risk classification under Article 6(3) (narrow procedural/refinement tasks) will still be required to register them in the EU database, reversing previous proposals that would have exempted them.
- **New Article 5 Prohibition ("Nudifiers" and CSAM):** The agreement bans AI systems designed to generate or manipulate non-consensual sexually explicit/intimate content or child sexual abuse material (CSAM). Providers of general-purpose image/video tools must actively prevent these foreseeable outcomes, subject to a transitional period ending December 2, 2026.
- **Relief for Small Mid-Cap Companies (SMCs):** Regulatory relaxations previously restricted to SMEs—such as simplified technical documentation, proportionate penalties, and less prescriptive quality management systems—are now extended to SMCs.
- **Industrial and Machinery Carve-out:** Embedded AI subject to the Machinery Regulation is removed from the direct application of the AI Act; safety measures will instead be managed via delegated acts under that regulation.
- **Reinforced EU AI Office Powers:** The AI Office is granted exclusive competence over AI systems built on General-Purpose AI (GPAI) models where the model and system are developed by the same provider. It is equipped with new enforcement tools, including powers to conduct investigations, perform on-site inspections, accept binding commitments, and issue fines.

### Verbatim Quotes

From the Gibson Dunn Client Alert:

"High-risk obligations for stand-alone Annex III systems are deferred to **2 December 2027**; for AI embedded in regulated products under Annex I, to **2 August 2028**." "The agreed text replaces the Commission's originally proposed conditional trigger mechanism with these fixed dates." "The agreement extends the existing legal basis for processing special-category personal data for bias detection and correction — previously limited to providers of high-risk systems — to all AI systems and general-purpose AI models, subject to a strict necessity standard..."

From the Travers Smith Briefing:

"Providers seeking exemption from high-risk classification for certain systems (e.g. systems limited to performing a narrow procedural task or refining the results of tasks previously completed by a human) will still need to register them in the EU database for high-risk systems, albeit with reduced information requirements. This represents a reversal from earlier drafts..."

"Processing special category data for the purpose of bias detection and correction returns to a strict necessity test, requiring providers to demonstrate that no less intrusive means exist to achieve the same objective."

### **Enterprise Risk Management Implications**

For corporate legal and risk teams, this agreement shifts the compliance timeline but does not relax the eventual standards:

1. **Prepare for August 2, 2026 Transparency:** Do not let the HRAIS delay distract from the fact that transparency and disclosure requirements under Article 50 are still active for August 2026.
2. **Review Bias-Correction Data Practices:** The rollback to a "strict necessity" test means enterprises cannot easily process special category data (like race or gender) for bias testing without proving that no alternative, less intrusive methods exist.
3. **Establish SMC Compliance Pathways:** Medium-sized enterprises should assess if they qualify as Small Mid-Caps to take advantage of the simplified documentation and proportionate penalty structures.

### **Sources**

- EU AI Act Omnibus Agreement — Postponed High-Risk Deadlines and Other Key Changes
- EU agrees to delay key AI Act compliance deadlines | Travers Smith

## **FCRA and ICRAA Class Action Against Eightfold AI: Reframing AI Recruiting Liability Around Consumer Reporting**

In a groundbreaking class action lawsuit, *Kistler et al. v. Eightfold AI Inc.* (Case No. 4:26-cv-01768, N.D. Cal. 2026), plaintiffs have bypassed traditional AI discrimination claims to target the privacy and procedural implications of automated hiring. The lawsuit alleges that Eightfold AI secretly compiles and scores job applicants on a 0-to-5 scale without providing the disclosures, authorizations, and dispute-resolution procedures required under the federal Fair Credit Reporting Act (FCRA) and California's Investigative Consumer Reporting Agencies Act (ICRAA).

The litigation has entered a critical phase, with Eightfold AI filing a formal Motion to Dismiss and the court establishing a timeline for a precedent-setting ruling in late summer 2026.

### **Procedural Posture and Briefing Schedule**

The case, originally filed in California state court in January 2026, was removed to the U.S. District Court for the Northern District of California. The docket reveals the immediate timeline:

- **April 20, 2026:** Eightfold AI filed its formal Motion to Dismiss (Filing 29).
- **June 18, 2026:** Under the current briefing schedule, the plaintiffs (Erin Kistler and Sruti Bhaumik) are set to file their opposition to the motion to dismiss.
- **August 4, 2026:** U.S. District Judge Yvonne Gonzalez Rogers is scheduled to hold a hearing on the Motion to Dismiss in Oakland, California.

### Core Legal Battlegrounds on the Motion to Dismiss

Legal analysts and commentators point to several statutory hurdles that Judge Gonzalez Rogers must resolve, which will define the boundaries of AI vendor liability under credit reporting laws:

1. **The "Consumer Reporting Agency" (CRA) Definition:** To be held liable, Eightfold must meet the statutory definition of a CRA—meaning it assembles or evaluates consumer information "for the purpose of providing consumer reports to third parties." Eightfold argues that it is merely a software provider operating on data intentionally shared by candidates or provided by employers.
2. **The "Transaction or Experience" Exemption:** Under the FCRA, reports containing information "solely as to transactions or experiences between the consumer and the person making the report" are exempt. The court must decide if AI-generated assessments derived from third-party data scraped across the web (such as social media profiles, location data, and online activity) fall outside this exemption.
3. **The "Consumer Report" Threshold:** The court must determine whether an AI-generated "Match Score" or "likelihood of success" assessment constitutes a "consumer report" when used to filter out candidates before any human review occurs.

### The Vendor "Liability Squeeze" and the Algorithmic Pincer

The *Eightfold* litigation represents a massive strategic shift for plaintiffs' attorneys. By framing automated screening as a procedural credit reporting issue rather than an algorithmic bias issue, plaintiffs bypass the high bar of proving discriminatory intent or disparate impact.

This creates a devastating "pincer movement" when paired with *Mobley v. Workday* (where the court held that an AI recruiting vendor can be sued as an "agent" of employers under Title VII/ADEA). While *Workday* attacks discriminatory outcomes, *Eightfold* attacks automated processes. For enterprise employers, this widens the "liability squeeze":

- **Contractual Gaps:** Most third-party AI vendor contracts disclaim compliance warranties and cap liability to minimal subscription fees.
- **Regulatory Isolation:** Although the Consumer Financial Protection Bureau (CFPB) rescinded its 2024 guidance claiming algorithmic scores are FCRA-covered in 2025, private plaintiffs are successfully using the statute directly to seek statutory damages (\$100 to \$1,000 per willful violation). With Eightfold's database covering over one billion profiles, the financial exposure for both vendors and their enterprise clients is immense.

## Sources

- AI Hiring Under Fire: What the Eightfold Lawsuit Means for Every Employer Using Algorithmic Screening
- AI Hiring Tools and Consumer Reports: Understanding the Eightfold Litigation
- Kistler et al v. Eightfold AI Inc. 4:2026cv01768 Docket

# May 28, 2026 Cycle Summary: The Procedural Battlelines of AI Hiring under the FCRA

This research cycle focused on the critical procedural and briefing developments in *Kistler et al. v. Eightfold AI Inc.* (Case No. 4:26-cv-01768, N.D. Cal.), a landmark class action that seeks to classify AI-driven candidate ranking and scoring platforms as "consumer reports" under the Fair Credit Reporting Act (FCRA) and California's Investigative Consumer Reporting Agencies Act (ICRAA).

Because the case bypasses traditional algorithmic bias/discrimination claims to focus on procedural transparency and consent, its outcome will establish a massive precedent for the entire AI recruiting and HR technology industry.

## Key Developments Surfaced This Cycle

1. **The Motion to Dismiss Has Been Filed:** On **April 20, 2026**, Eightfold AI filed its formal Motion to Dismiss (Filing 29), officially initiating the judicial battle over whether credit reporting statutes apply to AI candidate-screening tools.
2. **Precedent-Setting Schedule Set:** The Northern District of California has established a clear briefing and hearing timeline:
  - **June 18, 2026:** Plaintiffs' opposition brief is due.
  - **August 4, 2026:** A crucial hearing on the Motion to Dismiss will take place before U.S. District Judge Yvonne Gonzalez Rogers in Oakland, California.
3. **Core Statutory Intersections:** The upcoming judicial evaluation will center on whether Eightfold's platforms meet the statutory definition of a "Consumer Reporting Agency" (CRA), whether its "Match Scores" derived from scraped data constitute "consumer reports," and whether the tools fall under the FCRA's exemption for "transactions or experiences."
4. **The Algorithmic "Pincer Movement" and Enterprise Squeeze:** Together with *Mobley v. Workday* (reaffirming vendor agency liability under Title VII/ADEA), the *Eightfold* litigation forms a devastating pincer targeting automated hiring. While *Workday* attacks discriminatory outcomes, *Eightfold* attacks secretive automated processes. This forces enterprise employers to urgently re-paper vendor contracts, conduct independent audits, and demand transparency regarding data sources and scoring logic to close the widening gap between contractual risk allocation and expanding legal liability.

**Sources**

- FCRA and ICRAA Class Action Against Eightfold AI: Reframing AI Recruiting Liability Around Consumer Reporting
- AI Hiring Under Fire: What the Eightfold Lawsuit Means for Every Employer Using Algorithmic Screening
- AI Hiring Tools and Consumer Reports: Understanding the Eightfold Litigation
- Kistler et al v. Eightfold AI Inc. 4:2026cv01768 Docket