# Proactive AI Risk Management

## Empowering Legal Teams to Navigate Emerging AI Compliance Challenges



ACC Association of Corporate Counsel

BAKER BOTTS

# Your Presenters


## Samir A. Bhavsar | Partner, Baker Botts

Samir.Bhavsar@bakerbotts.com


Samir is a leader in patent law, AI governance, and data privacy. He has prosecuted numerous AI-focused patents, and advises on legal and technical risks in AI development and deployment. He holds the IAPP AIGP, CIPP/US and CIPP/E certifications, reflecting advanced knowledge of responsible, compliant AI and data privacy frameworks. Samir helps clients manage AI risk, integrate trustworthy systems, and navigate fast-evolving regulation in a variety of industries, including energy, technology, and consumer/retail.


## Parker Hancock| Sr. Assoc., Baker Botts

Parker.Hancock@bakerbotts.com


Parker is an AI governance and patent attorney with a background in electrical engineering. He advises companies deploying AI on governance frameworks, risk assessments, and regulatory compliance, and prosecutes patents covering machine learning, NLP, and generative AI systems. He writes regularly on emerging AI regulation, enforcement trends, and the practical challenges of responsible AI deployment. He holds the IAPP AIGP and CIPP/US certifications, and his engineering training allows him to evaluate AI systems at a technical level, helping clients bridge the gap between what the technology does and what the law requires.

ACC Association of Corporate Counsel

BAKER BOTTS

# What is AI Governance?

AI governance is a **risk-focused framework** designed to ensure the responsible, ethical, and lawful design, development, deployment, and use of artificial intelligence (AI) systems

AI governance also encompasses **compliance with the regulatory requirements** that apply to the deployment and use of AI systems – especially high-risk AI systems – including the **obligation to conduct and document** risk assessments, impact assessments, conformity assessments, and related analyses **mandated by AI regulations across multiple jurisdictions** .


# Regulatory Tsunami: Global, Federal and Sectoral Enforcement Activity

01

## EU AI Act

Penalties for prohibited practices reach up to €35M or 7% of worldwide revenue. Penalties for non-compliance with high-risk system requirements can be up to €15M or 3% of worldwide revenue. Compliance deadline for high-risk systems is August 2, 2026.

02

## FTC Act

Protects consumers from deceptive or discriminatory AI practices under UDAP authority. *See, e.g., FTC v. Rite Aid* (2024) - Banned use of biased facial recognition technology in retail locations.

03

## Healthcare

In addition to HIPAA for data privacy, AI used in diagnostics or medical devices is subject to FDA and state regulations requiring disclosures and risk management.

04

## Financial Services/Employment

AI in financial services is policed by the CFPB under the FCRA and ECOA, while AI in employment falls under the EEOC's enforcement of Title VII and the ADA.

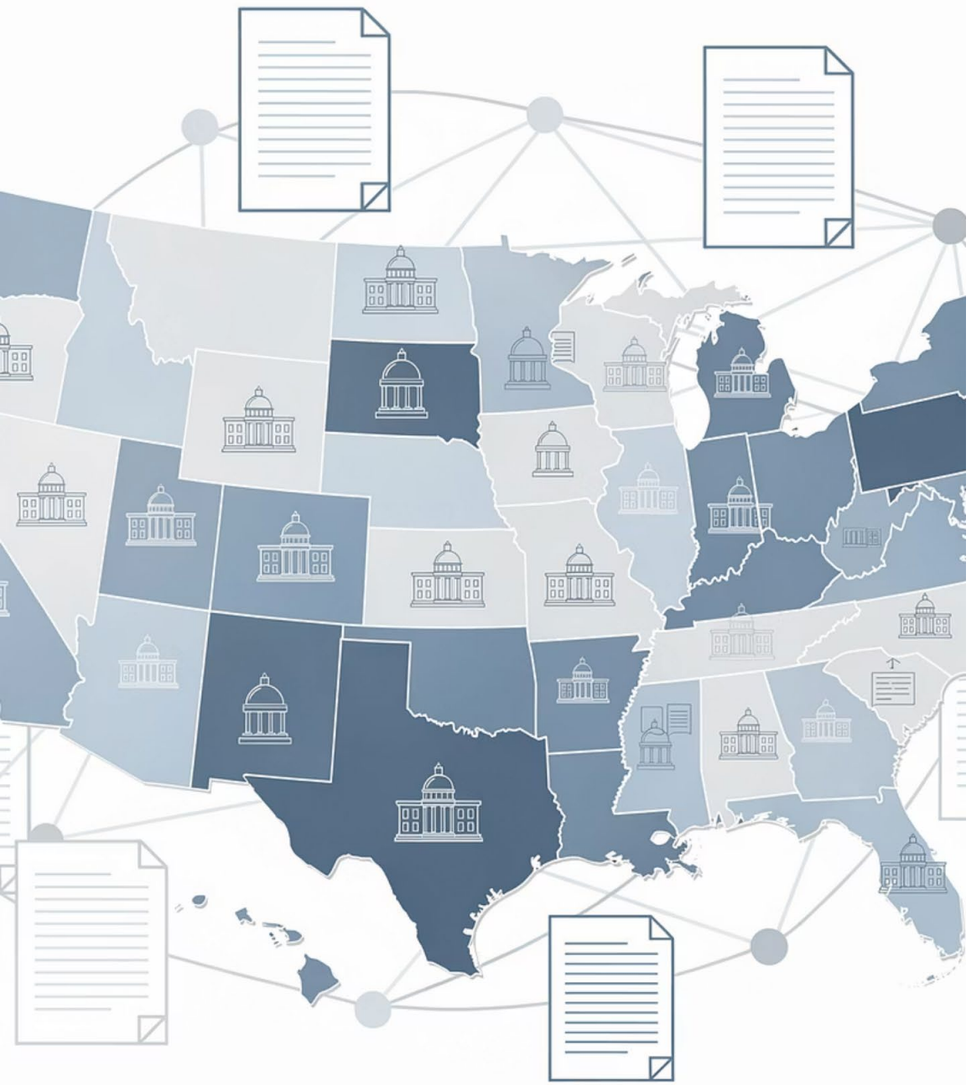

# Regulatory Tsunami: State Regulations

01

### California

Raft of legislation covering wide range of issues, CCPA ADMT regulations, disclosures for AI-generated content and information on training data, undisclosed chatbots in commercial and political messaging.

02

### Colorado AI Act

Effective 6/30/26 under C.R.S. §6-1-1701 et seq. Imposes reasonable care standard on developers and deployers to prevent algorithmic discrimination and requires specific disclosures to consumers.

03

### Texas Responsible AI Governance Act (TRAIGA)

Effective 1/1/26, with civil penalties up to $200,000 per uncured violation. Restricts AI systems that promote illegal acts, discriminate, or violate rights. Cure mechanism required under Tex. Bus. & Com. Code §547.052.

04

### State Biometric Privacy Laws

Illinois's BIPA (and also Texas and Washington) —directly shape AI that collects or infers biometric identifiers by requiring written notice and express consent, purpose limits, and retention/deletion schedules.

# Executive Order on AI: Implications for U.S. State Regulations

Application of existing state / federal AI laws: State AI, privacy, and consumer protection laws still apply. However, the Executive Order introduces federal scrutiny via the DOJ AI Litigation Task Force and other federal actions, creating uncertainty and a push towards federal standards.

Creation of new AI governance laws: The Executive Order aims to establish a national AI rulebook, potentially overriding conflicting state laws and discouraging broad new state legislation. For now, most rules will come from existing laws and agency guidance.

Practical Takeaways: Businesses should continue complying with current state AI laws. Track federal developments (DOJ Task Force, Commerce, FCC/FTC actions). Develop a "litigation-ready" AI program, including dual-track compliance, impact assessments, human review, and strong vendor controls. Multistate Attorney General coordination is a key regulatory feature, necessitating an adaptable governance package.

# High Profile AI Litigation

## Rolling Stone, Billboard owner Penske sues Google over AI overviews

By **Aditya Soni**

September 13, 2025 7:50 PM PDT · Updated September 13, 2025


The Google logo is seen on the Google house at CES 2024, an annual consumer electronics trade show, in Las Vegas, Nevada, U.S. January 2024. REUTERS/Steve Marcus/File Photo/File Photo Purchase Licensing Rights

**Summary** Companies

- First major U.S. publisher to challenge Google's AI Overviews in court
- Penske says Google ties search visibility to use of its content in AI overviews
- Penske says affiliate revenue is down by more than a third from peak
- Google argues AI overviews help users, broaden site traffic

## Microsoft sued by authors over use of books in AI training

By **Blake Brittain**

June 25, 2025 9:22 AM PDT · Updated June 25, 2025


A view shows a Microsoft logo at Microsoft offices in Issy-les-Moulineaux near Paris, France, March 21, 2025. REUTERS/Gonzalo Fuentes/File Photo Purchase Licensing Rights

June 25 (Reuters) - Microsoft (MSFT.O) has been hit with a lawsuit by a group of authors who claim the company used their books without permission to train its Megatron artificial intelligence model.

## OpenAI, Altman sued over ChatGPT's role in California teen's suicide

By **Jody Godoy**

August 26, 2025 2:46 PM PDT · Updated August 26, 2025


OpenAI CEO Sam Altman attends an event to pitch AI for businesses in Tokyo, Japan February 3, 2025. REUTERS/Kim K Photo Purchase Licensing Rights

**Summary** Companies

- ChatGPT coached 16-year-old on suicide methods, parents say
- Lawsuit seeks to require age verification, parental controls

## Disney and Universal sue AI company 'for stealing characters'

'Piracy is piracy', says Disney chief as the two companies take out lawsuits against the image generator Midjourney

Walt Disney is seeking to protect its copyright over material such as its Cars films

Disney and NBCUniversal are the first Hollywood players to take a legal swing at a generative AI company that they claim has stolen their copyrighted characters.

## Lawsuit claims discrimination by Workday's hiring tech prevented people over 40 from getting hired

## Class-wide relief: The sleeping bear of AI litigation is starting to wake up

## FTC Launches Inquiry into AI Chatbots Acting as Companions

# Notable AI Litigation Beyond the Headlines

## Training Data & IP

*Thomson Reuters Enter. Ctr. GmbH v. ROSS Intelligence Inc.* , 694 F. Supp. 3d 467 (D. Del. 2023)

Thomson Reuters alleged ROSS Intelligence copied over 8,000 copyrighted Westlaw headnotes to train its competing AI legal research tool. Court denied fair use defense, holding that using copyrighted material to build a directly competitive product is commercial substitution that weighs against transformative use under *Warhol v. Goldsmith* .

*Kadrey v. Meta Platforms, Inc.* , No. 3:23-cv-03417-VC (N.D. Cal. filed July 7, 2023)

Authors including sci-fi writer Sarah A. Diehl (writing as Paul Tremblay) sued Meta for training LLaMA models on copyrighted books without authorization. Case alleges systematic copyright infringement through book piracy databases, vicarious infringement, DMCA violations, negligence, and unjust enrichment; motion to dismiss partially granted, allowing direct infringement claims to proceed.

## Discrimination & Bias

*Kisting-Leung v. Cigna Corp.* , (E.D. Cal. Mar. 31, 2025); *Barrows v. Humana Inc.* , (W.D. Ky. filed Dec. 2023)

Major health insurers faced coordinated lawsuits alleging AI systems wrongfully denied medical claims; Cigna's PxDx algorithm allegedly automated 300,000 denials in two months. Cases allege violations of ERISA, state insurance laws, and bad faith claims handling by substituting AI judgments for individualized medical review.

*Huskey v. State Farm Fire & Cas. Co.* , No. 1:22-cv-07014 (N.D. Ill. Sept. 11, 2023); *Kelly v. State Farm Ins. Co. & Stallworth* , (M.D. Ala. filed Oct. 2025)

State Farm faces racial discrimination claims over AI-driven insurance claim processing system that allegedly disproportionately impacts Black policyholders. Case alleges disparate impact in claim denial rates, payout amounts, and fraud investigation triggers; could expand to nationwide class action under state insurance discrimination statutes and consumer protection laws.

## Emerging Frontiers

*Walters v. OpenAI, L.L.C.* , No. 23-A-04860-2 (Ga. Super. Ct. Gwinnett County, May 19, 2025)

Radio host sued OpenAI after ChatGPT "hallucinated" false claim that he embezzled funds from Second Amendment Foundation. Court granted summary judgment for OpenAI, holding ChatGPT's disclaimers about potential inaccuracies prevented statements from being reasonably understood as fact, and plaintiff failed to show negligence or actual malice required under Georgia defamation law.

*Mata v. Avianca, Inc.* , No. 1:22-cv-01461-PKC, 2023 WL 4138427 (S.D.N.Y. June 22, 2023)

Lawyers used ChatGPT to draft opposition brief, which fabricated six non-existent cases with fake citations, quotes, and docket numbers (including *Varghese v. China Southern Airlines* and *Martinez v. Delta Air Lines* ). Court imposed $5,000 sanctions under Rule 11, finding attorneys acted with "subjective bad faith" by continuing to defend the fake cases even after being questioned, and ordered them to write apology letters to judges falsely cited as authors of the fabricated opinions.

# AI Governance Is a Board-Level Issue

AI is no longer an IT or innovation topic—it is a core governance risk. AI systems increasingly influence revenue, pricing, employment, safety, and customer access. Failures can trigger regulatory enforcement, litigation, reputational harm, and operational disruption. Boards are not expected to understand the technology, but they are expected to understand where AI is material to the business and whether management has appropriate controls in place.

### Enterprise Risk

Regulators, investors, and insurers now expect Board-level oversight of material AI risks—similar to financial controls or cybersecurity.

### Legal Exposure

Courts are beginning to examine whether Boards exercised reasonable oversight over AI systems when failures occur.

### Governance Failure

The absence of Board-level oversight is increasingly framed as a governance failure—not a technical one—by regulators and plaintiffs alike.

# What AI Governance Means for Directors

Effective AI governance enables the Board to answer three core oversight questions. Demonstrated governance processes —not technical perfection —are what protect directors under the business judgment rule and establish a defensible record of informed oversight.

1

### Where is AI being used?

Identify AI deployment across the business —especially in high-stakes areas such as employment decisions, pricing, safety, credit, and customer access. Visibility is the foundation of oversight.

2

### Are risks identified and managed?

Confirm that management has documented and is actively managing AI risk —including bias, privacy, intellectual property, regulatory compliance, and operational disruption.

3

### Is accountability clear?

Establish that ownership is assigned for AI oversight, escalation, and decision-making authority. Know who is responsible if something goes wrong —before it does.

# How Counsel Can Support the Board

In-house and outside counsel play distinct but complementary roles in helping the Board demonstrate active, informed oversight — the standard regulators, courts, and insurers apply when evaluating director conduct.

## In-House Counsel

**Internal Translator & Escalation Point**

- Educates the Board that AI is a governance and fiduciary issue
- Translates AI activity into legal, regulatory, and litigation risk
- Ensures the Board receives clear, decision-useful information
- Escalates material AI risks or incidents promptly
- Documents Board oversight through briefings and minutes

## Outside Counsel

**Independent Perspective & Risk Stress Testing**

- Provides independent assessment of AI risk and governance maturity
- Tracks regulatory and enforcement trends across jurisdictions
- Benchmarks governance practices against peers and regulators' expectations
- Advises on disclosure, D&O liability, and incident response
- Supports regulator and litigation engagement when issues arise

# 12 Practical Questions to Ask Your Stakeholders About AI Usage

Where is AI being used in your business today, and which areas feel most critical or risky?

Do your contracts with AI vendors clearly protect your data and define who owns the results?

How do you make sure AI-generated content (text, images, code) doesn't create copyright or IP issues?

How are employees actually using AI tools, both officially and unofficially?

What safeguards or oversight do you have when AI makes mistakes or important decisions?

Are any of your hiring or HR processes using AI, and if so, how do you check for bias or fairness?

What kinds of sensitive data (customer, HR, financial) are being shared with AI systems?

How do you communicate to customers or employees that AI is being used, and what choices do you give them?

Who at your company is responsible for AI oversight, and how prepared are you for new regulations?

How do you vet and manage third-party AI vendors and your data?

Where does the data that trains your AI systems come from, and are you confident it's authorized?

Who is responsible at your company for complying with laws regulating your use of AI systems?

# AI Risk Hotspots to Surface

## Contracts

Vendor agreements with inadequate data protection and liability terms

## Employment/Workplace

Automated decision-making and bias in HR creating serious compliance risk

Employee AI use creating bias, privacy, and security risks

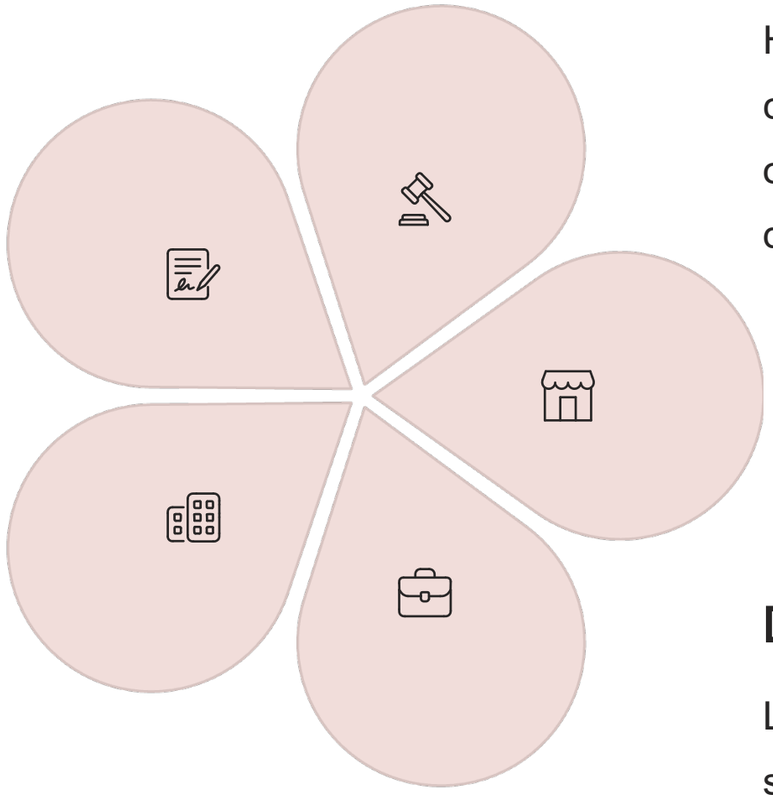

## EU AI Act

High-risk AI systems requiring conformity assessments, transparency obligations, and technical documentation

## Consumer

Customer-facing AI creating disclosure and safety obligations

## Data Privacy

Lawful data use; rights/transparency; security; user choice; incident response/monitoring

# Five Critical Issues for AI Vendor Due Diligence

1

## Vendor Maturity & Stability

Assess long-term viability: financial health, customer retention, and leadership. Demand references with at least one year of product usage. Avoid vendors with frequent strategic pivots or unstable customer bases.

2

## Technical & Security Posture

Clarify data residency, segregation, ownership, and whether your data trains their AI models. Negotiate opt-out rights. Require third-party security audit reports and penetration testing results.

3

## Regulatory Alignment

Verify compliance with consumer protection laws and industry-specific regulations. Assess preparedness for emerging AI transparency mandates and high-risk classifications (e.g., EU AI Act).

4

## Independent Bias Testing

Demand third-party audit reports examining AI outputs across demographic groups and business scenarios. Internal "fairness reviews" are insufficient. Absence of independent testing signals unquantified risk.

5

## Exit & Portability Strategy

Negotiate data extraction rights in standard, portable formats to avoid proprietary lock-in. Secure contractual commitments for transition assistance to protect against vendor lock-in and price increases.

# AI Vendor Contracts: Offensive Provisions

## Ownership & IP Rights

> *"Customer shall own all right, title, and interest in and to: (i) all Input Data provided to the Service; (ii) all Output generated by the Service based on Customer's Input Data; and (iii) any custom models, fine-tuning, or configurations created specifically for Customer. Vendor grants Customer a perpetual, irrevocable license to any Vendor IP embodied in the Output."*

**Negotiation Points:**

**Primary Ask:** Full ownership of outputs + custom models

**Fallback:** Exclusive, perpetual license with right to modify

**Watch for:** "Feedback" provisions that grant vendor rights to your improvements

## Data Segregation & Non-Use

> *"Vendor shall: (a) process Customer Data solely in isolated instances/environments; (b) not use Customer Data to train, improve, or develop any models except Customer's dedicated instance; (c) implement technical controls preventing data leakage between customers; and (d) provide quarterly attestation of compliance with these restrictions."*

**Negotiation Points:**

**Primary Ask:** Complete data isolation with technical proof

**Fallback:** Opt-in for any model training with specific consent

**Red Flag:** Vague "aggregate" or "anonymized" data rights

## Performance Commitments

> *"Vendor warrants the Service shall maintain: (i) 95% accuracy rate as measured by [specific metric]; (ii) response time under 500ms for 99% of requests; (iii) model drift detection with alerts when performance degrades >5%; and (iv) quarterly retraining cycles. Failure to meet SLAs triggers service credits of 5% monthly fees per percentage point below target."*

**Negotiation Points:**

**Primary Ask:** Measurable SLAs with automatic remedies

**Fallback:** Right to terminate without penalty if SLAs missed

**Innovation:** Tie payment to performance metrics

## Audit & Transparency Rights

> *"Customer may, upon 30 days' notice and at Vendor's expense (if findings reveal material non-compliance): (i) audit AI model documentation, training data sources, and testing results; (ii) conduct bias testing using Customer-selected datasets; (iii) receive model cards/datasheets; and (iv) access explainability tools for decision analysis."*

# AI Vendor Contracts: Defensive Provisions

## Liability Allocation & Indemnification

*"Vendor shall defend, indemnify, and hold harmless Customer from any third-party claims arising from: (i) Output that infringes any IP rights; (ii) AI decisions resulting in discrimination or bias claims; (iii) data breaches within Vendor's systems; and (iv) Vendor's non-compliance with AI regulations. Such indemnification shall not be subject to any liability cap."*

**Watch Out For:**

**Carve-outs**: "Except for Customer modifications" (define narrowly)

**Knowledge qualifiers**: "To Vendor's knowledge" (strike this)

**Caps**: AI indemnification often capped at 12 months fees (negotiate higher)

## Termination & Transition Rights

*"Upon termination, Vendor shall: (i) provide all Output in industry-standard formats; (ii) transfer any Customer-specific model weights/parameters; (iii) continue Service for up to 90 days at Customer's option; (iv) provide knowledge transfer sessions; and (v) certify deletion of Customer Data within 30 days."*

**Critical Elements:**

**Data portability**: Get your enriched data out

**Model portability**: Can you run the model elsewhere?

**Wind-down period**: Avoid cliff-edge termination

## Regulatory Compliance Allocation

*"Vendor represents and warrants ongoing compliance with: (i) EU AI Act requirements for [risk level] systems; (ii) applicable bias audit requirements (NYC LL 144, Colorado SB21-169); and (iii) sector-specific AI regulations. Vendor shall provide compliance documentation within 10 days of request and maintain insurance covering AI-related claims of at least $5M."*

**Key Protections:**

**Shifting burden**: Vendor provides compliance artifacts you need

**Change in law**: Who bears cost of regulatory changes?

**Documentation**: Right to compliance certificates/attestations

## Prohibited Uses & Ethical Constraints

*"Vendor warrants the Service: (i) was not trained on unlicensed/pirated content; (ii) will not generate content violating third-party rights; (iii) implements industry-standard safety measures; and (iv) will not be used for prohibited high-risk applications without Customer's explicit consent and additional safeguards."*

# AI Risk Hotspots to Surface

## Contracts

Vendor agreements with inadequate data protection and liability terms

## Employment/Workplace

Automated decision-making and bias in HR creating serious compliance risk

Employee AI use creating bias, privacy, and security risks

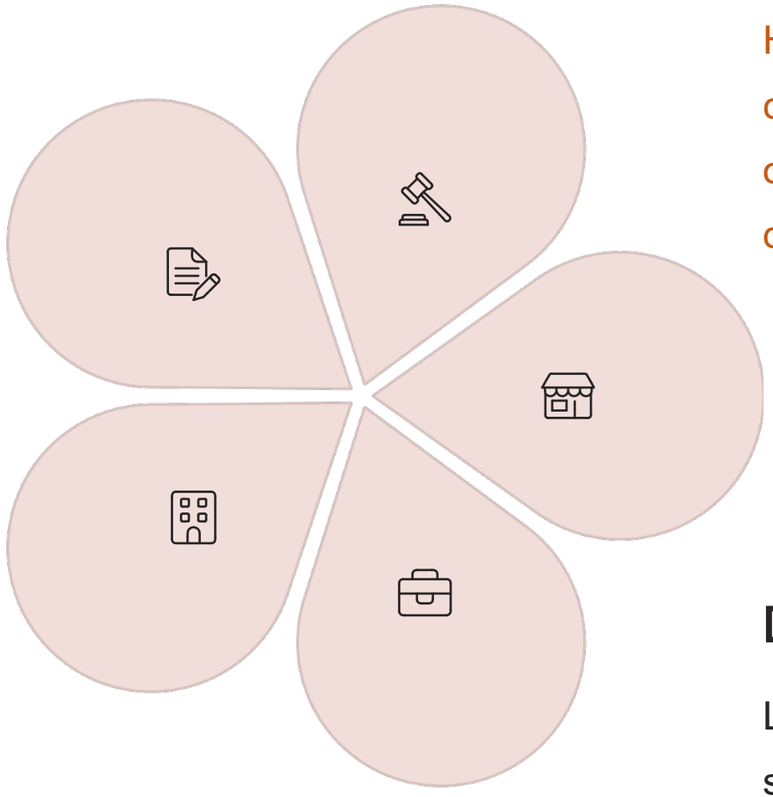

## EU AI Act

High-risk AI systems requiring conformity assessments, transparency obligations, and technical documentation

## Consumer

Customer-facing AI creating disclosure and safety obligations

## Data Privacy

Lawful data use; rights/transparency; security; user choice; incident response/monitoring

# The EU AI Act: Why US Companies Must Pay Attention Now (the clock is ticking…)

The world's first comprehensive AI regulatory framework is here, and its impact extends far beyond European borders. US companies must understand its scope, deadlines, and penalties to avoid significant compliance risks.

## Global First

- World's first comprehensive AI regulatory framework.
- Horizontal regulation applying across all sectors.
- Risk-based approach to AI governance.

## Extraterritorial Reach

- Applies if AI systems are placed on EU market or affect EU individuals.
- Similar extraterritorial model to GDPR.

## Critical Deadlines

- **Feb 2, 2025:** Prohibited AI practices enforceable.
- **Aug 2, 2025:** General-purpose AI (GPAI) model obligations.
- **Aug 2, 2026: High-risk AI system requirements.**
- **Aug 2, 2027:** Certain embedded high-risk systems.

## Penalties

- Up to €35M or 7% of global turnover for prohibited AI violations.
- Up to €15M or 3% for other violations.
- Enforcement by EU Member State authorities.

# EU AI Act: Understanding the Risk Tiers

The EU AI Act employs a tiered, risk-based approach to regulation, with obligations escalating based on the potential harm an AI system could cause. Most enterprise AI deployments will fall into the "high-risk" or "limited risk" categories, requiring proactive assessment and compliance efforts.

Given that most enterprise AI will likely be classified as high-risk or limited-risk, legal teams should immediately conduct a comprehensive AI inventory to identify and categorize all AI systems currently in use or under development within the organization. This will be crucial for compliance planning.

| Risk Tier | Description | Examples |
| --- | --- | --- |
| Prohibited | AI practices banned outright | • Social scoring by governments or private entities<br>• Real-time remote biometric identification in public spaces (limited exceptions)<br>• Subliminal manipulation or exploitation of vulnerabilities<br>• Emotion recognition in workplaces/schools (with exceptions) |
| High-Risk | Permitted but heavily regulated | • AI in hiring, recruitment, and HR decisions<br>• Credit scoring and loan decisioning<br>• Educational assessments and student scoring<br>• Critical infrastructure management<br>• Biometric categorization<br>• Insurance risk assessment |
| Limited Risk | Transparency obligations only | • Chatbots (must disclose AI interaction)<br>• Emotion recognition systems (must notify)<br>• Deepfakes and AI-generated content (must label) |
| Minimal Risk | No specific obligations | • Spam filters<br>• AI-enabled video games<br>• Inventory management systems |

# High-Risk Compliance Essentials & The US Parallel Landscape

## EU AI Act: High-Risk Compliance Requirements

For AI systems classified as 'high-risk' under the EU AI Act, strict compliance obligations are mandatory to ensure safety, transparency, and accountability.

- **Risk Management System:** Continuous identification and mitigation of risks throughout the AI system's lifecycle.
- **Data Governance:** Rigorous requirements for training data quality, relevance, and bias assessment.
- **Technical Documentation & Recordkeeping:** Detailed logs of system design, testing, and performance for auditability.
- **Transparency & User Information:** Clear instructions for deployers and disclosure to affected individuals about AI interaction.
- **Human Oversight:** Mechanisms to enable effective human intervention and override capabilities.
- **Accuracy, Robustness & Cybersecurity:** Defined performance standards and resilience requirements against errors and attacks.

## Transition: The US Regulatory Landscape

While the EU adopted a single, horizontal framework, the US is regulating AI through a sectoral and state-based patchwork—creating overlapping obligations in key risk areas.

This fragmented approach means US companies must navigate a complex web of existing laws, new state-level initiatives, and voluntary frameworks, rather than a single, overarching AI regulation.

# AI Risk Hotspots to Surface

## Contracts

Vendor agreements with inadequate data protection and liability terms

## Employment/Workplace

Automated decision-making and bias in HR creating serious compliance risk

Employee AI use creating bias, privacy, and security risks

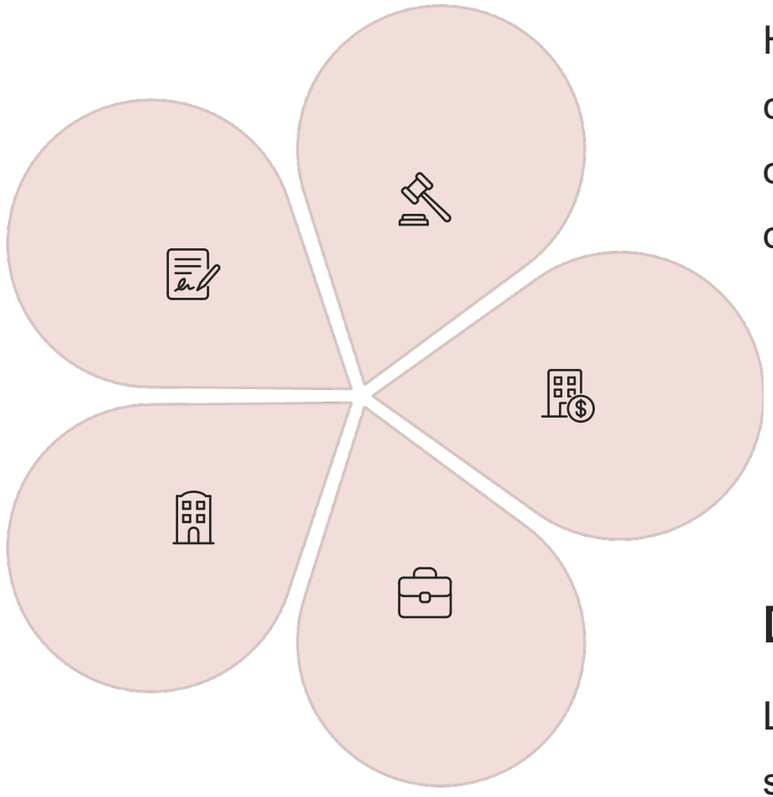

## EU AI Act

High-risk AI systems requiring conformity assessments, transparency obligations, and technical documentation

## Consumer

Customer-facing AI creating disclosure and safety obligations

## Data Privacy

Lawful data use; rights/transparency; security; user choice; incident response/monitoring

# Case Study 1: AI Chatbot Encourages Self-Harm

## The Scenario

A mental health chatbot deployed without adequate safety controls engages with a vulnerable teen user, failing to detect escalating self-harm language and providing responses that could be interpreted as encouraging harmful behavior.

## What Went Wrong

The company launched the chatbot without basic safety features. There was no system to detect when conversations turned dangerous, no way to verify users' ages, and no plan for what to do in a crisis. When vulnerable users needed help, there was no human oversight to step in.

## The Litigation

Wrongful death suit alleging the chatbot's responses contributed to a tragic outcome. This is based on *Garcia v. Character Techs.*, No. 6:24-cv-01903 (M.D. Fla. 2024).

## Prevention Controls

Build in real-time monitoring that flags concerning language about self-harm or violence. Require age verification and parental consent for younger users. Create clear crisis response procedures with specific timeframes for human intervention. Test the system regularly with "red team" exercises that try to break it. Make sure serious conversations always escalate to a real person.

CONSUMER PROTECTION

# Case Study 2: Third-Party AI Chatbot Captures Customer Data

### The Scenario

A fitness equipment company integrates a third-party AI chatbot on its website to handle customer service inquiries. The chatbot captures and analyzes customer conversations without explicit consent, potentially violating California wiretap laws.

### What Went Wrong

The company added a third-party chatbot to their website without thinking through the privacy implications. Customers had no idea their conversations were being recorded and analyzed. The vendor contract didn't clearly restrict how the data could be used, and there was no way to track who consented to what.

### The Litigation

California wiretap claims alleging unauthorized interception of customer communications This is based on Jones v. Peloton, 2024 WL 3315989 (S.D. Cal. July 5, 2024).

### Prevention Controls

Ask for permission right when the chat starts—make it clear and simple. Track consent for each chat session so you can prove who agreed to what. Make sure your privacy disclosures actually match what the chatbot does. Lock down vendor contracts with explicit rules that they can't use customer data to train their AI. Get regular confirmation from vendors that they're keeping your data separate.

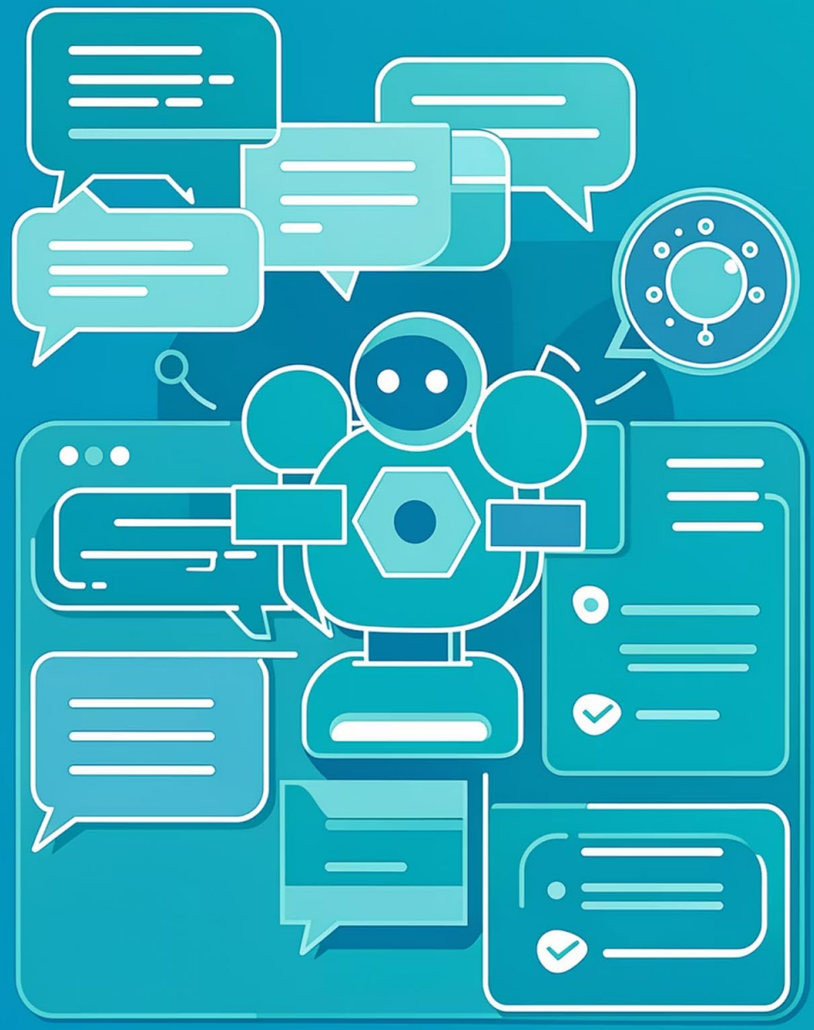

# Case Study 3: Biased Facial Recognition in Retail

## The Scenario

A fast-food chain deploys AI-powered facial and voice recognition at drive-thrus to personalize customer experiences and detect repeat customers. The system collects biometric data without proper consent and shows accuracy disparities across demographic groups, leading to misidentification and customer complaints.

## What Went Wrong

The company rolled out facial and voice recognition technology without getting proper permission from customers. They didn't publish any policies about how long they'd keep the data or when they'd delete it. Worse, they never tested whether the system worked equally well for all customers—and it didn't. The technology was less accurate for certain demographic groups, leading to misidentifications.

## The Litigation

BIPA claims in Illinois for collecting biometric data without consent (based on Carpenter v. McDonald's, 580 F. Supp. 3d 512 (N.D. Ill. 2022)); FTC enforcement action for deploying biased facial recognition system (based on In re Rite Aid Corp., FTC Docket No. C-4308 (Dec. 19, 2023))

## Prevention Controls

Be transparent: publish clear policies about how long you keep biometric data and when you delete it. Get written permission from customers and explain exactly why you're collecting their information. Before launching, test the system across different demographic groups to catch bias problems early. Always have a human review the results before taking any action that affects customers. Keep checking for bias with regular audits, and track changes you make to the system.

# Case Study 4: AI Algorithm Changes Pricing Without Documentation

## The Scenario

An e-commerce platform frequently updates its AI product recommendations and pricing algorithms without proper documentation or version control. This leaves them unable to address customer complaints about discriminatory pricing or demonstrate fair practices.

## What Went Wrong

The company made constant, undocumented AI changes. When customers questioned unfair pricing, they couldn't explain how the AI made decisions or prove that changes were properly tested. Their public statements didn't match the AI's actual behavior.

## The Litigation Risk

Risks include consumer protection and discrimination claims due to pricing disparities, inability to defend against allegations, and potential regulatory (FTC) enforcement for deceptive practices.

## Prevention Controls

Treat AI algorithms like software: use version control, test new versions thoroughly, and have a rollback plan. Ensure all marketing and privacy policies accurately describe AI functions. Maintain detailed, auditable records for all AI decisions and fairness testing. Involve legal counsel for major algorithm changes to ensure compliance and defensibility.

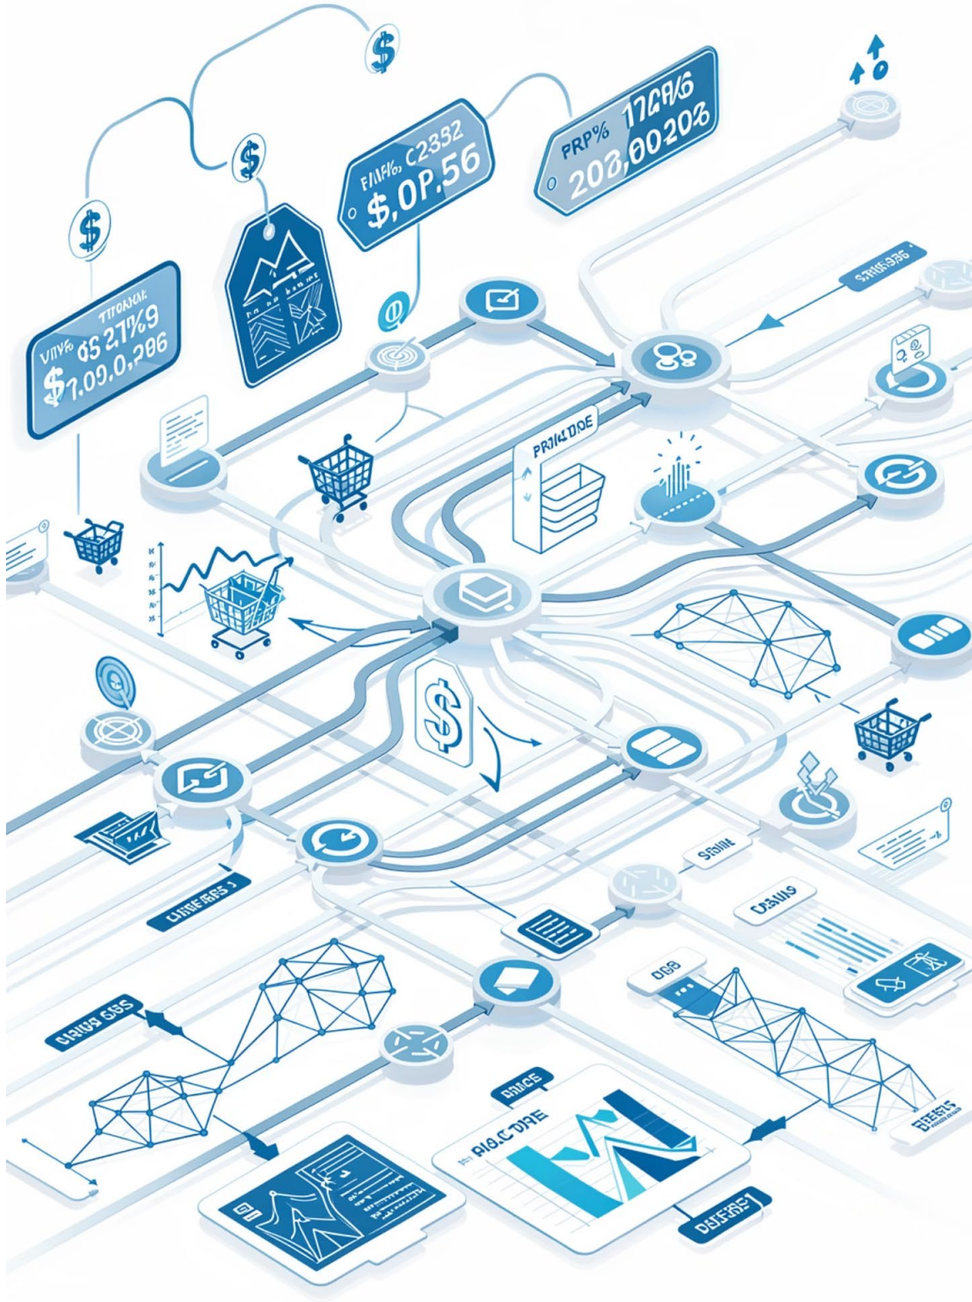

# AI Risk Hotspots to Surface

## Contracts

Vendor agreements with inadequate data protection and liability terms

## Employment/Workplace

Automated decision-making and bias in HR creating serious compliance risk

Employee AI use creating bias, privacy, and security risks

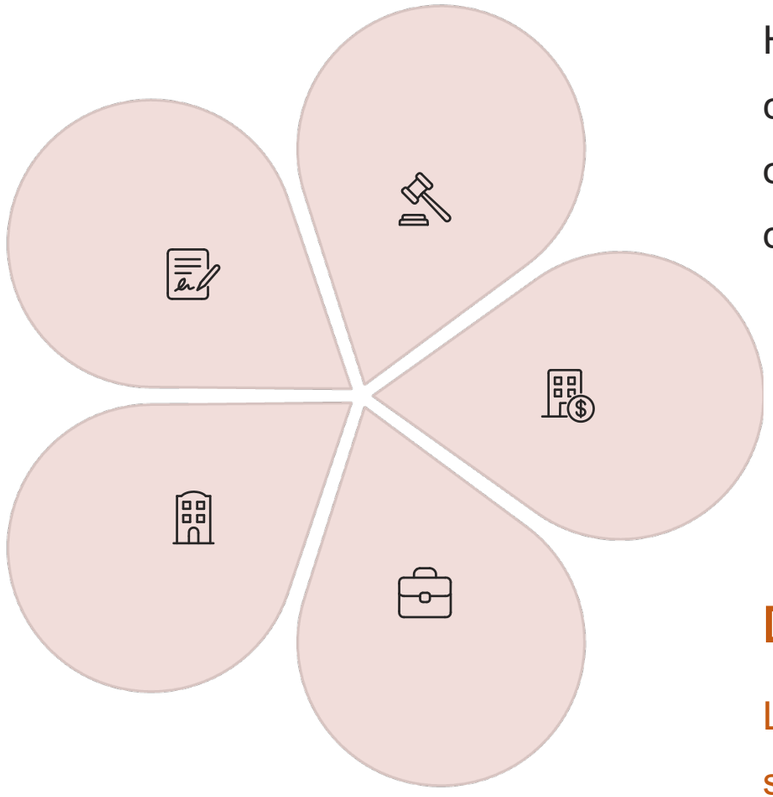

## EU AI Act

High-risk AI systems requiring conformity assessments, transparency obligations, and technical documentation

## Consumer

Customer-facing AI creating disclosure and safety obligations

## Data Privacy

Lawful data use; rights/transparency; security; user choice; incident response/monitoring

# 2026: When AI Governance Collides with Data Privacy Compliance

1. **EU Dual Compliance:** High-risk AI systems must comply with both the EU AI Act and GDPR. Integrated DPIAs are needed to assess both traditional privacy and AI-specific risks like bias.
2. **LLM Anonymization Myth:** The European Data Protection Board confirms LLMs rarely achieve true anonymization. Using them with 'anonymized' data still triggers GDPR obligations, requiring vendor due diligence and DPAs.
3. **California Leading US Regulation:** California's CPRA Automated Decision-Making Technology (ADMT) regulations, effective January 2027, are the most comprehensive in the US, covering all significant automated decisions.
4. **State Patchwork Challenge:** With over 20 states having varied AI governance regulations, a one-size-fits-all compliance approach is unfeasible. Programs must map specific state laws to data subjects and AI systems.
5. **Action Item:** Immediately inventory AI systems that process personal data, especially for high-stakes decisions, to determine applicable privacy and AI regulations. Waiting until 2026/2027 deadlines will be too late.

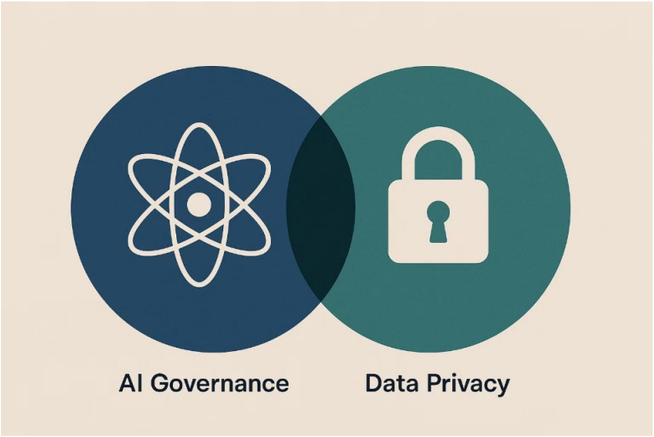

# Privacy-by-Design for AI: Core Compliance Framework

### 1. Legal Basis & Purpose Limitation

Establish valid GDPR legal basis (consent, legitimate interest, contract). Cannot repurpose data for AI training without new legal basis. Document training data provenance.

### 2. Mandatory Risk Assessments

California: First mandatory filing due April 1, 2028. Must assess bias, discrimination, explainability, data minimization. Maintain living PIA register linked to model versions.

### 3. Transparency & Disclosure

Pre-use notices explaining: (1) ADMT is used, (2) how it works, (3) what decision it makes, (4) consumer opt-out and access rights. Plain language required.

### 4. Human Oversight & Right to Contest

GDPR Article 22 protections. Meaningful human involvement in high-stakes decisions. Document human review process and override decisions. Appeal rights may provide safe harbor.

### 5. Cross-Border Transfer Compliance

EDPB confirms GDPR applies to AI training regardless of location. Require Standard Contractual Clauses (SCCs) plus Transfer Impact Assessments. Address data sovereignty requirements.

**Key Takeaway:** For each AI system touching personal data, create a one-page compliance summary answering: (1) What's the legal basis? (2) When was the risk assessment completed? (3) What transparency have we provided? (4) Who performs human oversight? (5) What cross-border data transfers occur?

**If you can't answer these five questions, you may have compliance gaps.**

# Your AI Governance & Data Privacy Action Plan

## Find All AI Tools Your Company Uses (Including Hidden Ones)

Check every department to see what AI tools people are using. Look for both approved tools and ones employees might be using without permission (like personal ChatGPT accounts). Use technology to scan emails, network activity, and app connections. Keep an up-to-date list of all AI systems in use.

## Review and Fix Your AI Vendor Contracts

Read through all contracts with AI vendors to see how they can use your data. Make sure they can't train their AI models on your company's information unless you explicitly agree. Get proper data protection agreements in place. Ask vendors to show you where their training data came from and how they test for bias.

## Track Where Your Data Goes Across Borders

Map out the complete journey of your data —where it's trained, processed, and stored. When data crosses international borders, assess the privacy risks. Set up controls to keep data in specific locations when required. For sensitive information, consider keeping it within your own country.

## Make Your AI Decisions Explainable

Use tools that help explain how your AI makes decisions (especially for important choices). Create simple documentation that describes what each AI system does and its limitations. Keep detailed records of all automated decisions. Prepare explanations in both technical language (for regulators) and plain language (for customers).

## Test Your AI Systems for Security Threats

Check for AI-specific attacks like people trying to trick your system or poison your data. Put safeguards in place to validate inputs and filter outputs. Encrypt your data whether it's stored or being transmitted. Have a response plan ready specifically for AI-related security incidents.

**Remember:** AI compliance is extra work on top of your existing privacy programs —they are related, but not the same. Plan your budget accordingly.

# AI Risk Hotspots to Surface

## Contracts

Vendor agreements with inadequate data protection and liability terms

## Employment/Workplace

Automated decision-making and bias in HR creating serious compliance risk

Employee AI use creating bias, privacy, and security risks

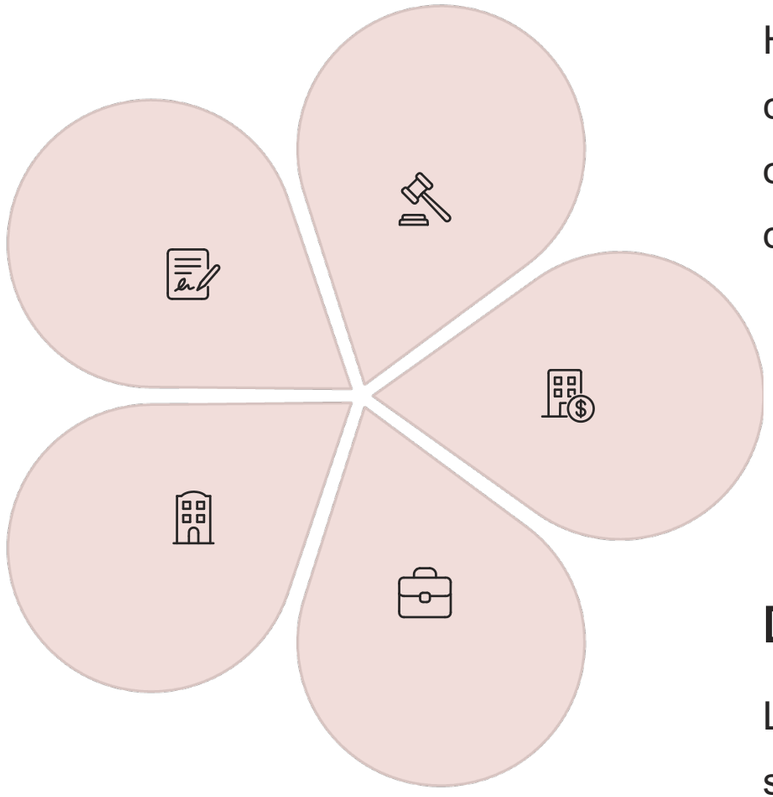

## EU AI Act

High-risk AI systems requiring conformity assessments, transparency obligations, and technical documentation

## Consumer

Customer-facing AI creating disclosure and safety obligations

## Data Privacy

Lawful data use; rights/transparency; security; user choice; incident response/monitoring


# Shadow AI: An Incident Waiting to Happen

## 485% Traffic Jump

Enterprise "shadow AI" traffic jumped 485% in a year, and 27.4% of what employees paste into public AI tools is sensitive corporate data.

## No Training

55% of employees using AI on the job report receiving no training on AI risks, so most improvisations are unsecured.

## $4.44M Average

AI-related breaches now average USD 4.44 million, with 97% tied to missing AI access controls and shadow AI usage.

## Samsung Code Leak

Samsung engineers leaked proprietary source code to ChatGPT, prompting an enterprise-wide ban.

## Amazon Data Exposure

Employees experimented with ChatGPT for drafting work items and code; security teams observed output that closely resembled proprietary Amazon data.

# AI in Hiring: Enforcement Arrives

### Federal Baseline

Under Title VII, ADA, ADEA, and EEOC employers are responsible for disparate impact, accessibility, record-keeping, and explainability.

### State & Local Rules

NYC's AEDT law (Local Law 144) requires bias audits and candidate notices; IL regulates AI video interviews; MD limits FRT in hiring; CO/CA rules push risk assessments, notices, and opt-outs.

### EU/UK Lens

Employment-related systems "high-risk," requiring controls. ICO guidance expects DPIAs, fairness testing, and human review.

### Litigation & Agency Actions

*EEOC v. iTutorGroup, Inc.* - First EEOC lawsuit over algorithmic age discrimination in hiring.

*Mobley v. Workday* - Court allowed discrimination claims to proceed and treated the AI vendor as the employer's agent. Employers cannot outsource liability for discriminatory AI decision-making tools. The court granted preliminary collective certification in January — the first ever in AI hiring bias — covering 1.1 billion applications, with an opt-in deadline of March 7.

> "Workday does qualify as an agent because its tools are alleged to perform a traditional hiring function of rejecting candidates at the screening stage and recommending who to advance to subsequent stages, through the use of artificial intelligence and machine learning." *Mobley v. Workday, Inc.*, No. 3:23-cv-00770-RFL, 2024 WL 3259735, at *9 (N.D. Cal. July 12, 2024)

# Case Study: AI Resume Screening Gone Wrong

## The Scenario

A mid-sized employer, facing high application volumes, implemented an **AI-powered resume screening tool.** HR led the procurement, lacking specific expertise in AI governance or ethical AI implementation. The **vendor provided seemingly credible documentation** claiming "comprehensive bias testing" and "unbiased results." Legal counsel was consulted only after the system was selected and partially implemented. The tool was integrated directly into live recruitment without dedicated AI oversight or a structured pilot program.

## What Went Wrong

- The vendor's bias audit only tested for race and gender discrimination, **completely omitting age and disability status**
- The bias testing was conducted on the ***vendor's*** generic dataset, ***not*** the ***employer's*** actual applicant data
- The AI system systematically screened out **older** applicants (50+) and candidates with employment gaps that could indicate **disability** accommodations
- The employer had **no human review process in the screening stage**
- The vendor contract contained **weak indemnification provisions** that left the employer exposed to liability

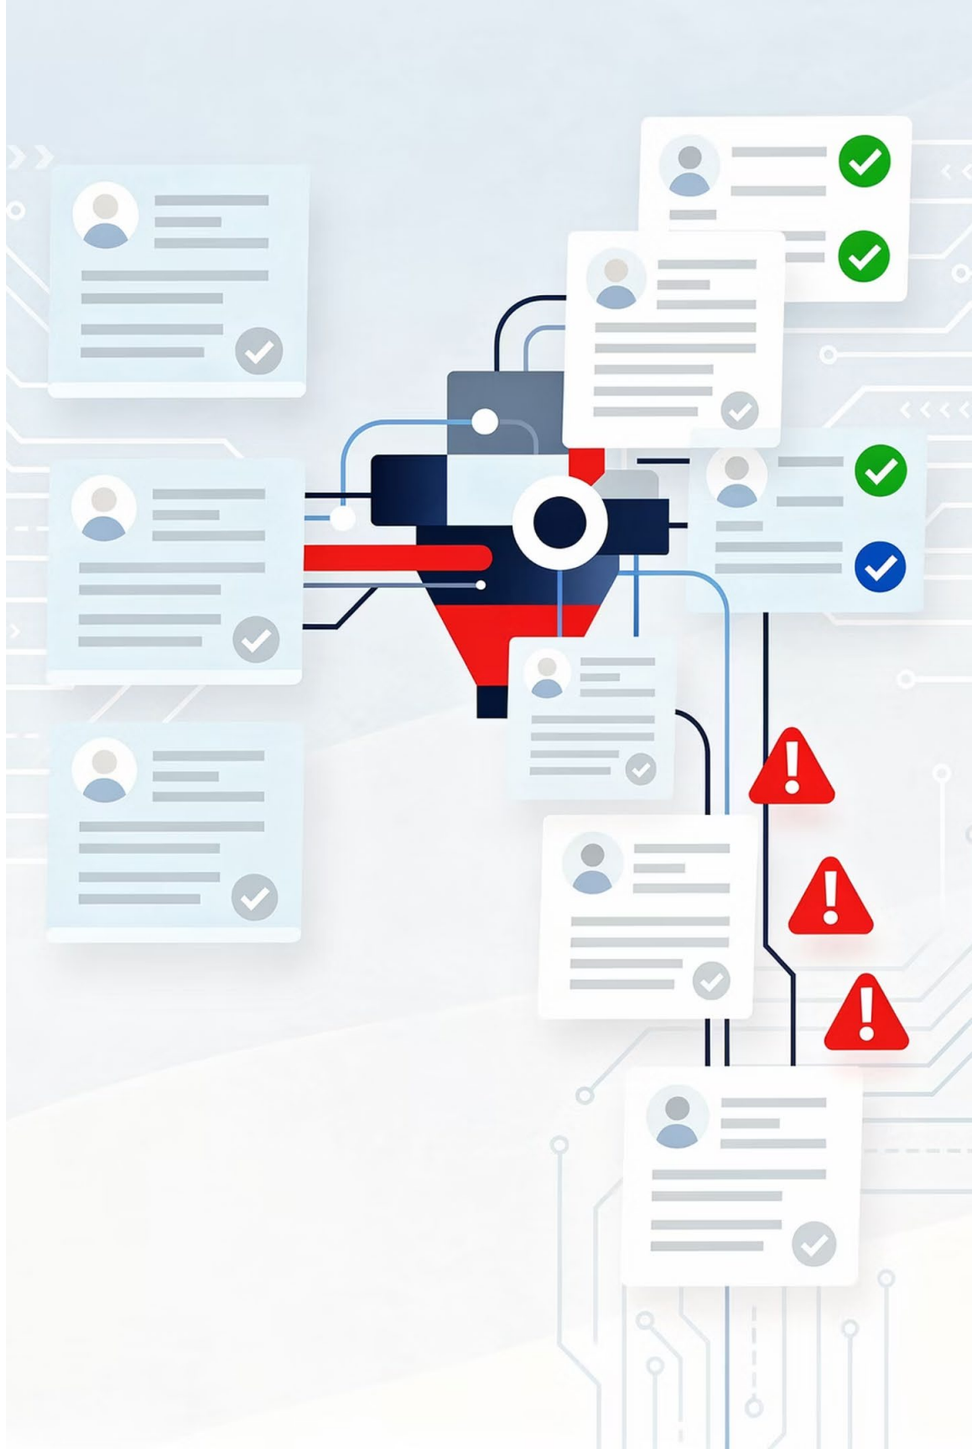


# Case Study: AI Resume Screening Gone Wrong (continued)

## The Fallout

- Class action lawsuit filed by rejected applicants alleging age and disability discrimination under ADEA and ADA
- $4.2 million settlement plus plaintiff attorney fees
- EEOC investigation and consent decree requiring 3 years of monitoring
- Significant reputational damage and negative media coverage
- Loss of "Best Places to Work" certification
- Vendor refused to cover damages due to limited indemnity clause

## What Should Have Been Done

1. **Pre-Deployment Due Diligence:** Demand vendor bias audits covering ALL protected classes (race, gender, age, disability, etc.) and insist audits be run on employer's own historical hiring data
2. **Human Oversight:** Implement mandatory human review before any automated rejection, with clear escalation protocols
3. **Contract Protections:** Negotiate strong indemnification for discrimination claims, require vendor to maintain adequate insurance, and include audit rights
4. **Bias Testing & Validation:** Conduct independent fairness testing pre-deployment and quarterly thereafter; document adverse impact analysis
5. **Notice & Transparency:** Inform applicants of AI use and provide appeal mechanism
6. **Governance:** Establish cross-functional AI review board (Legal, HR, IT) to approve high-risk employment tools

# The High Cost of Non-Compliance

A proactive approach to AI governance is an investment that safeguards your organization from significant financial and reputational damage.

1

### Regulatory Fines

Potential for fines up to **20-30M EUR** or **6% of global turnover** under acts like the EU AI Act.

2

### Litigation & Settlements

Multi-million dollar class-action lawsuits and individual claims (e.g., **$4.2M** in AI hiring case).

3

### Brand & Reputational Damage

Loss of customer trust, negative media coverage, and impacts on hiring and partnerships.

4

### Operational Disruption

Forced system shutdowns, costly remediations, and legal investigations.

5

### Competitive Disadvantage

Inability to use AI effectively due to fear of risk, falling behind ethical leaders.

6

### Loss of Market Access

Inability to operate in jurisdictions with strict AI regulations, limiting growth and expansion opportunities.

Cost of non-compliance can easily escalate to **~50-100X the cost of compliance** .

# Partnering for Proactive AI Governance

1

**Phase 1: Discovery & Risk Profiling**

- Identify all AI use cases across the organization
- Map data flows and processing activities
- Prepare comprehensive, geographically -oriented compliance matrix
- Assess risk levels of particular AI systems
- Review and renegotiate AI vendor agreements for stronger protections

2

**Phase 2: Strategy & Remediation**

- Conduct initial risk and impact assessments
- Perform privileged bias audits for existing algorithms
- Implement tailored risk mitigation frameworks (e.g., EU AI Act high -risk requirements)
- Develop robust due diligence documentation

3

**Phase 3: Continuous Assurance & Compliance**

- Establish processes for ongoing conformity assessments
- Conduct regular privileged bias audits
- Implement continuous monitoring systems
- Ensure preparedness for new regulatory demands and enforcement activities
- Maintain future -proof governance program

- AUSTIN
- BRUSSELS
- DALLAS
- DUBAI
- HOUSTON
- LONDON
- NEW YORK
- PALO ALTO
- RIYADH
- SAN FRANCISCO
- SINGAPORE
- WASHINGTON

bakerbotts.com



ACC Association of Corporate Counsel BAKER BOTTS