AI Compliance in the Workplace: What You're Really Up Against

As more companies (and their employees) adopt AI, they are unwittingly opening themselves up for legal issues that most business leaders never saw coming. AI is everywhere now—drafting contracts, analyzing customer data, automating financial decisions, creating marketing content, and making recommendations that affect everything from loan approvals to insurance claims. Each use case brings potential liability for biased outcomes, privacy violations, intellectual property infringement, and regulatory non-compliance that traditional risk management frameworks weren't designed to handle.

The regulatory response is scrambling to catch up, with federal agencies applying 1960s civil rights concepts to 21st-century machine learning while states and cities write new AI laws faster than most legal teams can track them. But here's what we've learned after helping dozens of companies navigate AI deployment: the solution isn't avoiding AI—it's deploying it thoughtfully, with compliance built in from the start rather than retrofitted after problems emerge.

DUNLAP
BENNETT&
LUDWIG

AI Compliance in the Workplace: Scope and Accountability

AI is now part of everyday work decisions. Hiring screens, performance analytics, productivity copilots, and security monitoring systems all use algorithms to influence decisions that affect people's livelihoods. And many hiring decisions are now made through a host of new platforms that employe large language models as the backbone for automated decision-making systems (ADS) in recruiting. That's why AI compliance in the workplace isn't just an IT task—it spans human resources, legal, privacy, security, procurement, and business leadership. The accountability model that works in 2025 looks like this: **the business owns the decision, humans remain in the loop, and vendors don't get to absorb all of your liability by contract.**

Over the past two years, employee AI usage has jumped and shadow use has grown with it. Surveys show many employees use AI tools weekly while a large share do so without approval, sometimes pasting sensitive data into public models. The signal here is simple: policy without enforcement doesn't hold. Clear rules, named owners, role-specific training, and technical controls like data loss prevention matter as much as the legal fine print.

Think of a common scenario. A recruiter opens a resume-ranking tool and sees a top-10 slate with confidence scores. If the tool nudges choices toward patterns in historic data, risk rises. Accountability means documenting the tool's purpose, inputs, outputs, and limits, testing outcomes for disparate impact, and giving recruiters a clear escalation path when something "looks off." As people often say when a dashboard gets too slick, "Pretty charts don't equal good judgment."

Federal Oversight and Existing Employment Laws

Despite policy swings in Washington, employers remain fully on the hook under existing civil rights and labor laws. Federal agencies have reinforced that algorithmic tools don't bypass Title VII or the ADA, and that employers are responsible for outcomes even when a third-party vendor builds the model. The broader federal climate also includes active privacy and consumer protection enforcement that touches model training data, disclosures, and deceptive claims about AI capabilities.

EEOC Title VII and ADA Obligations

Title VII prohibits disparate treatment and disparate impact based on protected characteristics. The ADA adds duty-to-accommodate considerations that AI can easily frustrate if not designed with accessibility in mind. It is important to note, however, that on January 27, 2025, the EEOC removed AI-related guidance from its website following President Trump's executive order rescinding Biden's AI executive order. However, the EEOC's Strategic Enforcement Plan for Fiscal Years 2024–2028, which prioritizes regulating employers' use of technology including AI and machine learning through enforcement, remains in effect until a quorum of commissioners can modify or revoke it. We've seen litigation trends where courts allow claims to proceed when plaintiffs plausibly allege algorithmic screening harmed protected groups. The cases that stick tend to have one thing in common: **employers who couldn't show they'd done any bias testing or impact analysis before deployment.**

Department of Labor Worker Well-Being Principles

The Department of Labor has articulated worker well-being principles for AI that center human oversight, transparency, explainability, and the right to contest consequential decisions. However, the DOL has noted that its "AI & Inclusive Hiring Framework" and "Artificial Intelligence Best Practices" guidance may be outdated or not reflective of current policies. While guidance language has shifted at the federal level, the operational bar in 2025 is clear: explain what the tool does, train managers on appropriate use, and keep a human decision-maker accountable for employment outcomes.

Agency Enforcement Trends and Guidance

Enforcement is converging from multiple directions. The FTC has probed AI companies about training data and risk mitigation, signaling broader scrutiny of claims and safety practices. Financial regulators have flagged deepfakes and model risks in cybersecurity programs. Expect cross-agency collaborations when AI creates discrimination, privacy, or deception risks that land in overlapping jurisdictions.

State and Local Laws Employers Must Track

States and cities are filling the vacuum with detailed obligations. A practical takeaway for multi-state employers: track jurisdictional triggers, build to the strictest common denominator, and document exceptions where a local rule goes beyond the baseline.

California Employment Regulations Regarding Automated-Decision Systems and California Civil Code §§ 1798.100 et seq., § 7001 et seq.

California employers now face overlapping requirements from both the Civil Rights Council and the Privacy Protection Agency

The California Civil Rights Council's Employment Regulations Regarding Automated-Decision Systems became effective on October 1, 2025. These regulations apply to employers with five or more employees in California and explicitly extend FEHA's anti-discrimination protections to automated decision systems. The new law requires a number of items specifically related to AI data. Employers must keep employment records, including ADS data, for a minimum of four years (increased from two years) and keep records of anti-bias testing of ADS.

On July 24, 2025, the California Privacy Protection Agency finalized regulations under the CCPA addressing automated decision-making technology, approved by the Office of Administrative Law on September 22, 2025. These create dual obligations where:

1. Employers using ADS have until January 1, 2027, to comply with the notice requirements, and
2. Businesses must conduct risk assessments for activities including the use of ADS for significant decisions in employment contexts.

Beginning in 2027, there will be new notice, access, and opt-out rights for individuals regarding the use of ADS in making significant decisions.

New York City Local Law 144

NYC's automated employment decision tools law requires annual independent bias audits, public posting of audit summaries, and candidate notices when AI substantially assists hiring or promotion decisions. The obligation is triggered by tools that play a predominant role in decisions, so employers have attempted to structure "human-predominant" workflows. Auditors and plaintiffs will look to substance over labels, so audit quality and clear documentation matter.

We've handled cases where companies thought they were safe because a human "reviewed" every AI recommendation. The problem? The human spent an average of twelve seconds per review and overrode the AI less than 2% of the time. **That's not meaningful human oversight—that's rubber-stamping.**

Colorado AI Act

Effective February 1, 2026, Colorado's law applies to high-risk AI used for consequential employment decisions. Covered employers must adopt risk management programs, run annual impact assessments, notify individuals when AI meaningfully influences decisions, and provide appeal mechanisms. The standard of "reasonable care" to avoid algorithmic discrimination will push vendors and employers to share testing data and correction rights.

Connecticut: S.B. 2, though currently stalled, is expected to be reintroduced in 2025 and would prohibit algorithmic discrimination with requirements similar to Colorado's standard

Illinois HB 3773

Effective January 1, 2026, Illinois amends its Human Rights Act to prohibit AI use that results in discrimination, requires worker notice across employment decisions, and bars proxies like ZIP codes that can serve as stand-ins for protected traits. Expect Illinois agencies to evaluate how employers validate and monitor third-party tools for adverse impact and accommodation compatibility.

Texas: House Bill 1709

The Texas Responsible Artificial Intelligence Governance Act, has been introduced with a potential effective date of September 2025.

Massachusetts

The proposed Artificial Intelligence Accountability and Consumer Protection Act (HD 396) would regulate high-risk AI systems with requirements for risk management programs and impact assessments.

Recent Court Decisions Affecting Vendor Liability for AI

On May 16, 2025, Judge Rita Lin in the Northern District of California granted preliminary certification for a nationwide collective action under the Age Discrimination in Employment Act (ADEA), allowing the lawsuit against the AI HR and finance platform Workday to proceed on behalf of all individuals aged 40 and over who applied through Workday's platform since September 2020, holding that Workday could potentially be held liable as an "agent" of employers, expanding potential liability for AI vendors. The Mobley ruling suggests that AI vendors could be directly liable for discrimination if their algorithm, acting as a delegated hiring function, unlawfully screens out protected groups.

The Legal Checklist for AI Deployment at Work

Use the checklist below as a working sequence. Treat it like a living control set you revisit as models evolve, data drifts, and local rules update.

Before Deployment: Risk Assessments and Impact Testing

Start with a comprehensive inventory. Map every system where AI influences employment decisions—hiring, promotion, termination, scheduling, performance evaluation, and compensation. Classify "consequential" uses that could significantly affect someone's job prospects or working conditions.

Run pre-deployment bias and impact tests. This isn't optional testing—it's liability insurance. Document data sources, representativeness, missing data issues, and model limitations. Record justifications for features and thresholds before you deploy, not after someone complains.

Complete data protection impact reviews for privacy laws. Identify sensitive attributes, role-based access controls, retention schedules, and cross-border data flows. Most AI employment systems process more personal data than companies realize.

Secure legal review against Title VII, ADA, and applicable state laws. Define acceptable error rates and remediation triggers before go-live. **The time to have these conversations is before deployment, when you still have options.**

During Deployment: Worker Notice, Human Oversight, and Safeguards

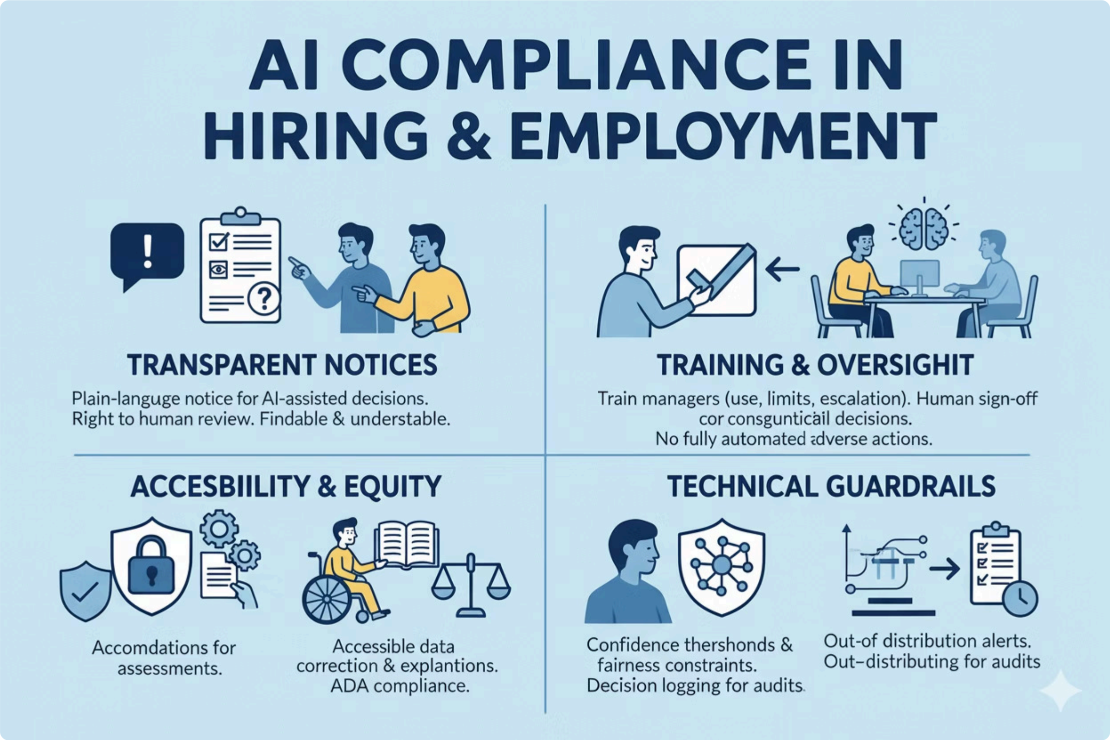
AI COMPLIANCE IN HIRING & EMPLOYMENT
TRANSPARENT NOTICES
Plain-languge notice for AI-assisted decisions. Right to human review. Findable & understable.
TRAINING & OVERSIGHIT
Train managers (use, limits, escalation). Human sign-off cor consgunticail decisions. No fully automated adverse actions.
ACCESBIILITY & EQUITY
Accomdations for assessments.
Accessible data correction & explantions. ADA compliance.
TECHNICAL GUARDRAILS
Confidence thershonds & fairness constraints. Decision logging for audits.
Out-of distribution alerts. Out–distributing for audits

Provide plain-language notices to applicants and employees when AI substantially assists decisions. Include information about the right to request human review where required by law. Make the notices findable and understandable—burying key facts in privacy policies doesn't satisfy notice requirements.

Train managers on appropriate use, limitations, and escalation paths. Require human sign-off for consequential decisions. Prohibit fully automated adverse actions. We've seen too many cases where "human oversight" meant a manager clicking "approve" on a screen they didn't understand.

Enable accommodations and alternatives for assessments that could screen out people with disabilities. Maintain accessible processes to correct data and explain outcomes. The ADA doesn't have an exception for algorithms.

Implement technical guardrails: confidence thresholds, fairness constraints, and alerts for out-of-distribution inputs. Log decisions and rationale for audits. When regulators investigate, they want to see the decision trail.

After Deployment: Monitoring, Audits, and Remediation

Monitor for drift and disparate impact on a defined cadence. AI systems can develop bias over time as they learn from new data or as underlying populations change. Compare cohorts over time, investigate root causes, and document adjustments.

Schedule independent audits where laws or risk levels demand it. Publish summaries when required by local rules like NYC Local Law 144. Budget for these audits—they're not cheap, but they're cheaper than class action settlements.

Provide appeal channels and timely remediation. Track cycle time from complaint to fix. Report material incidents to leadership and, where required, regulators. The appeals process needs to be real, not performative.

Data Privacy and Security Controls for Employee Data

AI runs on data, and employment data is sensitive by definition. Privacy and security controls do more than satisfy statutes—they build the worker trust that keeps AI programs viable.

01

Data Minimization, Retention, and Access Controls

Collect only the attributes needed for the stated purpose. Avoid proxy features like ZIP code that correlate with protected traits. We've seen hiring algorithms that used commute distance as a factor, not realizing it effectively screened out candidates from certain neighborhoods.

Set retention schedules specific to each AI use. Time-box training sets and purge derived artifacts that no longer serve a lawful purpose. Employment data that sits around indefinitely becomes a liability target.

Use least-privilege access and strong authentication for AI platforms. Segment model training environments from production HR systems. The people building AI models shouldn't have access to live employee records unless absolutely necessary.

02

Privacy Notices and Consent Where Required

Provide clear notices describing data types, sources, uses, and sharing. In opt-out or consent jurisdictions, honor worker choices and document preference handling. Don't make privacy choices theoretical—build them into system workflows.

Explain automated decision-making in understandable terms. Share key factors that influenced outcomes where law requires or fairness warrants. "The algorithm decided" isn't an explanation.

03

Incident Response for AI-Related Data Breaches

Extend incident response plans to AI-specific scenarios: model inversion, prompt injection, data poisoning, and deepfake risks. These aren't theoretical threats—we've seen all of them in real environments.

Practice tabletop exercises with HR, legal, PR, and security teams. Define clock-start for breach notification when model artifacts expose personal data. AI incidents can be more complex than traditional data breaches.

Vendor Due Diligence and Contract Requirements

Most employment AI arrives through vendors. Contracts are your primary control surface. Treat diligence as non-negotiable and verification as ongoing, not a one-time checkbox.

Bias Audit, Transparency, and Testing Rights

Require documented bias and impact testing, including methodology and cohort breakdowns relevant to your workforce. Don't accept vendor claims about "bias-free" algorithms without seeing the actual testing data.

Secure audit rights, model cards, and change logs. For black-box systems, negotiate outcome-level testing rights and performance guarantees. If you can't understand how the system makes decisions, you can't defend those decisions.

Data Use, Confidentiality, and IP Protections

Prohibit vendor training on your employee data without express approval. Define permitted uses, de-identification standards, and deletion timelines. Your employee data shouldn't become part of the vendor's next product release.

Protect prompts, outputs, and embeddings as confidential information. Clarify who owns fine-tuned models and derived works. These ownership questions matter more than most companies realize.

Indemnities, Warranties, and Liability Allocation

Seek indemnities for IP infringement, data breaches, and unlawful discrimination caused by the tool. Calibrate caps to the risk profile, not just fees. A vendor charging $50,000 annually shouldn't have a $50,000 liability cap for discrimination claims.

Include warranties for compliance with applicable AI, employment, and privacy laws, plus SLAs for remediation and incident response. Make the vendor's compliance obligations specific and measurable.

Governance, Training, and Ongoing Monitoring

Good governance turns policy into practice and practice into habit. Without it, tools drift, people improvise, and risk creeps back in.

AI Policy, Acceptable Use, and Manager Training

Publish an AI acceptable use policy that lists approved tools, banned inputs like sensitive personal data, and escalation contacts. Make it role-specific and plain-language. Generic policies don't change behavior.

Deliver training for HR, recruiters, managers, and IT that covers bias concepts, ADA accommodations, explanations, and practical do's and don'ts. The training should be specific to the tools they'll actually use.

Recordkeeping, Internal Audits, and Reporting

Maintain a system register with metadata: purpose, data sources, model owner, last audit date, and jurisdictions in scope. Keep decisions and explanations for at least the retention period tied to the employment action.

Report key risk indicators to leadership quarterly. Include adverse impact trends, audit findings, appeal volumes, and remediation times. Make the metrics meaningful, not just comprehensive.

Align AI Governance with ESG Goals and Board Oversight

Connect AI fairness and worker well-being metrics to ESG reporting. Include workforce impact, training access, and transparency measures in sustainability disclosures. Investors and stakeholders increasingly expect algorithmic accountability as part of responsible business practices.

Give the board visibility into high-risk uses and incident learnings. Assign a committee to oversee responsible AI alongside cybersecurity. Board members should understand the compliance risks and business implications of AI deployment.

FAQs and The Bottom Line

What are the rules for using AI in the workplace?

There's no single federal AI law yet. Instead, employers must comply with existing anti-discrimination laws like Title VII and the ADA, plus a growing set of state and local rules including the California laws, NYC Local Law 144, the Colorado AI Act, and Illinois HB 3773. Notice requirements, bias testing, human oversight, and appeal rights are recurring themes across jurisdictions.

The basic principle is consistent: AI can't be used to discriminate, and employers remain responsible for the decisions their systems make or influence.

What are the compliance concerns with AI?

The biggest concerns are disparate impact in hiring and promotion, privacy and security risks from sensitive employee data, lack of transparency in decision-making, inadequate accessibility for people with disabilities, vendor "black boxes" that can't be audited, and shadow AI use by employees without proper controls.

Regulators focus on bias, deceptive claims about AI capabilities, data governance, security controls, and traceability. The enforcement trend shows agencies coordinating across traditional silos.

What is the compliance standard for AI?

Standards vary by jurisdiction, but a practical baseline in 2025 includes reasonable care to prevent algorithmic discrimination, documented risk management, bias and impact testing, meaningful human oversight, clear notices to affected individuals, and auditable records.

Some jurisdictions codify specific requirements like impact assessments and appeal rights for consequential decisions. The trend is toward more detailed obligations, not less.

How is AI being used in compliance?

Leading programs apply AI to monitoring, investigations, policy management, and control testing, but they pair these tools with human review and robust audit trails. Organizations with mature governance use AI across risk and compliance functions while maintaining accountability and clear escalation paths.

The key is using AI to enhance compliance capabilities while ensuring the compliance AI itself meets regulatory standards.

The Bottom Line

AI compliance in the workplace is manageable, but it requires proactive planning rather than reactive fixes. The companies that integrate compliance into AI deployment from the beginning move faster and face fewer risks than those who try to retrofit compliance later.

The regulatory landscape will continue evolving, but the core principles are stable: prevent discrimination, maintain human oversight, be transparent about AI use, protect employee data, and monitor outcomes continuously. Get those fundamentals right, and you'll be positioned to adapt as specific requirements change.

If you're planning AI deployment in employment contexts, start with a comprehensive compliance assessment. The investment in upfront planning pays off in smoother implementation, lower legal risk, and better business outcomes.

Dunlap Bennett & Ludwig is a veteran-owned law firm with offices across multiple states. Our employment and privacy attorneys help companies navigate AI compliance while maintaining focus on their core business objectives.

To learn more about **Dunlap Bennett & Ludwig** and how we can help you, call today at **888-306-4030** or email us at **clientservices@dbllawyers.com**

DUNLAP BENNETT & LUDWIG