

Littler

Agentic AI in the Workplace Report 2026

NOVEMBER 2026

Littler

Agentic AI in the Workplace

NOVEMBER 3, 2026

AUTHORS:

Zoe Argento, Jan-Ove Becker, Michelle Devlin, Hannah Drury, Kimberly Duplechain, Christine Fitzgerald, Stephanie Goutos, Laura Hayward, RyAnn Hooper, Valerie Jackson, Allan King, Maureen Lavery, Tim Lee, Johanna Liedtke, Deborah Margolis, Devjani Mishra, Brooke Niedecken, Jim Paretti, Rebecca Pasternak, Niloy Ray, Joy Rosenquist, Barbara Russell, Ivie Serioux, Kathryn Siegel, Michael Skidgel, Britney Torres, Grace Waddell, Alice Wang, Michael Whitbread

Contents

Executive Summary 3

I. **Introduction: What Is Agentic AI and How Employers Are Using It** 4
- A. Agentic AI Compared to Precursors 4
- B. Distinguishing Features of Agentic AI and the Autonomy Spectrum 5
- C. Common Workplace Applications of Agentic AI 7
- D. Why Adoption Is Accelerating and What Makes It Durable 9

II. **The Current Regulatory Environment for Workplace AI** 12
- A. The U.S. Federal Landscape 12
- B. U.S. State and Local Regulation 14
- C. Global Framework 18

III. **Reconciling the Current Regulatory Environment with the Advances in Autonomous AI Systems** 22
- A. Mismatch: AI Employment Laws and Agentic Systems 22
- B. The "No Robo Bosses" Era: Employment Law in the Age of Agentic AI 24

IV. **Key Labor and Employment Law Considerations for Agentic AI** 27
- A. Legal Risks of Precursor AI 27
- B. Agency Principles and the Allocation of Workplace Authority 29
- C. Dynamism and Discrimination Risk 32
- D. The Ability to Access Data Across Systems and Data Privacy 33
- E. Optimizing Productivity, Labor Budget, and Wage and Hour Compliance 35
- F. How Agentic AI Shifts Labor-Management Relations 36
- G. How Agentic Autonomy Shapes Workforce Restructuring and Reductions in Force 39
- H. Leave and Accommodation with Agentic AI 41
- I. Autonomy, Access, and Trade Secret Protection 43
- J. Vendors, Third Parties, and Accountability 44

V. **The Litigation and Investigation Environment** 47
- A. Trends in Employment-Related AI Litigation 47
- B. Theories of Liability Likely to Extend to Agentic Systems 48
- C. Discovery Challenges Unique to Agentic Systems 49
- D. How Plaintiffs May Adapt Existing Employment Claims to Agentic AI 51

VI. **Key Takeaways** 53
- A. A Proactive Approach to Agentic AI: Decide, Design, Defend 54

Conclusion 56

Appendix A: Compendium of U.S. State Laws and International Laws Governing Use of AI 63

Executive Summary

Agentic artificial intelligence (AI) moves workplace AI from recommendation to action. Agentic AI systems are designed to do more than draft, score, and recommend. They can break a goal into steps, call on other software, and complete tasks with limited human direction, such as scheduling interviews, answering employee questions, routing requests, summarizing performance information, or drafting and distributing documents. An agent's ability to carry a process across multiple steps may reduce handoffs and repetitive work, with the potential to gain significant efficiencies.

The adoption curve for agentic AI is aggressive. Gartner's 2026 CIO and Technology Executive Survey found that AI agent deployment is expected to increase from 17% in 2026 to more than 60% within the next two years,[1] with at least 15% of day-to-day work decisions being made autonomously through agentic AI.[2]

However, plans for agentic AI may be moving faster than the frameworks for overseeing it. Seventy-four percent of organizations expect to use AI moderately in the next two years,[3] but only 13% think they have the right governance in place.[4] It is then not surprising that over 40% of agentic AI projects are predicted to be canceled by the end of 2027, citing cost, unclear business value, and inadequate risk controls.[5]

This Report focuses on a path for employers to capture the value of agentic AI by treating autonomy as a deliberate design decision, reimagining functions based on new capabilities, adapting governance to reflect emerging laws and developing legal doctrine, and preserving the record.

This Report offers an analytical framework for overcoming uncertainty and making strategic decisions with a legal foundation as follows:

- What agentic AI is, how it differs from earlier tools, and where employers are using it;
- How federal, state, local, and global rules treat workplace AI, and where they assume there is a human decision-maker;
- The mismatch between the current regulatory environment and advances in autonomous AI systems;
- How agentic autonomy changes risk in discrimination, agency and authority, data access and privacy, wage and hour, labor relations, restructuring, leave and accommodation, trade secrets, and with respect to vendors; and
- How litigation and discovery are developing, and how plaintiffs may adapt existing claims.

The Report concludes with a practical framework for AI-enabled workforce transformation based on defensible governance: Decide, Design, Defend.

I. Introduction: What Is Agentic AI and How Employers Are Using It

Agentic AI refers to an AI system designed not merely to generate content or recommendations, but to pursue defined goals by conducting independent reasoning, planning multistep work, using digital tools and data, taking actions, and adapting its approach as conditions change.[6] Simple AI agents are not new technology: familiar examples include antivirus software and navigation apps.[7] Agentic AI is a newer concept and can enable end-to-end workflows for key organizational processes. Compared with generative AI, which typically responds to a user prompt, agentic AI systems can formulate and revise plans, interact with software and other aspects of their environment, and execute tasks with limited or no human supervision.[8] In the workplace, that distinction is significant: an agent may move beyond drafting an email or summarizing a policy to initiating and coordinating workflow steps across enterprise systems, thereby expanding potential productivity gains as well as the need for clearly allocated human oversight, authority, accountability, and risk controls.

A. Agentic AI Compared to Precursors

Agentic AI builds on earlier forms of artificial intelligence but differs in the scope of the tasks it can perform, the degree of autonomy it can exercise, and the ways in which it interacts with people and enterprise systems.[9] Understanding the evolution from predictive systems that identify patterns and forecast outcomes, to generative systems that create content in response to prompts, to agents that can plan and execute multistep work, provides essential context for evaluating agentic AI's workplace applications, benefits, and risks.[10]

Predictive AI uses statistical models and machine-learning techniques to analyze historical and current data, identify patterns, and estimate the likelihood of a future event or outcome.[11] Unlike agentic AI, however, the predictive system ordinarily produces a forecast or recommendation for a person or another system to act upon rather than independently planning and executing the steps needed to take action.[12] In the workplace, an employer might use a predictive AI tool to forecast staffing needs by evaluating historical workload, seasonal demand, employee availability, and other operational data, enabling managers to plan schedules or recruitment needs more effectively. Agentic AI relies on predictive AI at almost every layer to predict which outcome is likely to lead to the best outcome.

Generative AI refers to AI systems that learn patterns from large datasets and use those patterns to produce new content such as text, images, audio, video, or software code in response to a user's prompt or other input.[13] A large language model (LLM), which typically forms the underpinning of commonly used text-based generative AI tools, is a type of generative AI model trained on very large collections of text to learn patterns and relationships in language and to generate text by predicting likely sequences of words or tokens in response to an input.[14] A conventional generative AI system produces requested content in response to a prompt but does not independently determine and carry out the subsequent steps needed to complete a broader business process.[15] For example, a human resources employee might

use a generative AI tool to create a first draft of a job description based on specified duties, qualifications, and organizational requirements, after which the employee would need to review and revise the output for accuracy, consistency, and legal compliance before using the job description to post a job vacancy. Agentic AI relies on generative AI as its core reasoning and output engine.

Natural language processing (NLP) is a type of AI that enables computer systems to analyze, interpret, and produce human language. It is not a separate category of AI.[16] The terms predictive and generative AI describe what a system does with prompts and data. NLP describes the language-related techniques and capabilities the system may use to process the data. In employment, NLP may support predictive AI tools, for example by classifying the topic and tone or sentiment of employees' responses to a workplace climate survey. NLP may also be used in generative AI systems used at work, such as a tool for recording a meeting and providing a transcript of the meeting to the attendees. NLP is a component of almost every stage of an agentic AI loop.

Computer vision is a type of AI that enables computer systems to extract and interpret information from visual inputs such as digital images, video, and scanned documents.[17] Like NLP, computer vision is considered a processing field or capability but is focused on visual rather than linguistic information, not as a category parallel to predictive or generative AI. A common workplace application for computer vision is a facial recognition system employees use to access a secure worksite.[18] Computer vision becomes a component of agentic AI whenever the agent has to perceive something that only exists visually.

B. Distinguishing Features of Agentic AI and the Autonomy Spectrum

Agentic AI is less a new kind of software than a new allocation of authority. The practical issue is not whether a system is "agentic," or even what it is capable of doing, but what it is permitted to do before a human weighs in. Autonomy is not intelligence; it is permission.[19] Two employers can deploy the same capability and create very different legal profiles: one lets a system only recommend candidates, while another lets it reject candidates and schedule interviews without case-by-case approval. Three features make that allocation of authority consequential.

1. Three Distinguishing Features

Multi-step decision-making.[20] Where predictive tools return a score and generative tools return content, an agent is given a goal and works out the steps to reach it. Asked to "produce three qualified finalists by Friday," a recruiting agent may revise a posting, select job boards, screen and rank applicants, schedule interviews, and draft communications, cycling through planning, acting, observing, and adjusting. While employment decisions are not necessarily occurring at every step, one or more consequential choices, such as a screening rule, data source, or threshold applied, may still occur well before a polished output reaches a person.

Human role. Agentic systems shift the human role from deciding each step to setting goals, defining guardrails, and handling exceptions. A person may be *in the loop* (an action cannot take effect until the person acts), *on the loop* (the system acts while the person monitors and can intervene), or *out of the loop* for a defined task. Those labels, however, do not by themselves establish whether human involvement is legally sufficient. Whether human involvement is substantive depends on design: whether the reviewer has enough information to understand the output, is equipped to evaluate it rather than default to it, and has authority to change the result.[21] Employers can reserve human judgment for actions that materially affect applicants or employees, using exception-based oversight for administrative or reversible steps.

Tool use and integration. Earlier AI tools produced content that a person then acted on. Agents act directly, connecting to applicant tracking, HRIS, scheduling, payroll, and email systems, moving from "here is a proposed offer letter" to "the offer was sent." Which systems may the agent access, and which previously siloed data may it combine? Which actions require escalation? Can an action be undone, and can its consequences? An agent's authority is defined by what its permissions allow, not by what the task requires, and the two can diverge. Permissions set what the agent can do, logs

show what it did, and recovery mechanisms determine what can be undone. Those limits matter most when one agent delegates to another and each handoff can make it harder to reconstruct where, when, and why a consequential choice was made.

2. The Autonomy Spectrum

Autonomy is better understood as a spectrum than as a binary characteristic. Agentic systems range from tools that require human approval before acting to systems capable of independently making decisions and executing them. One framework developed through the Cloud Security Alliance (CSA) describes six levels of autonomy along this spectrum, reflecting progressively decreasing human involvement.[22] For employment purposes, the decisive question is what role a person plays at the moment an action affects an applicant or employee? Four deployment modes capture the range:

●○○○

Recommend.

The agent analyzes and proposes; a person decides and acts (*e.g.*, the system suggests a shortlist and a recruiter decides whom to advance).

●●○○

Draft.

The agent prepares work product or stages an action; a person reviews, revises, and sends it (*e.g.*, interview questions, candidate correspondence, or performance reviews).

●●●○

Execute with approval.

The agent queues actions, but each consequential step waits at a human checkpoint (*e.g.*, rejection notices go out only on approval).

●●●●

Execute autonomously.

The agent acts end to end within defined boundaries while a person monitors and handles exceptions (*e.g.*, the agent schedules interviews or advances candidates under preset criteria).

Autonomy is a deployment choice, not a fixed product attribute. One deployment might autonomously schedule interviews, draft correspondence for review, and only recommend hiring decisions, so autonomy is best assessed action by action. It should also be reassessed as integrations, permissions, and vendor features change, because the relevant question is what the system is authorized to do now, not what the implementation plan said on launch day.

3. How Legal Exposure Shifts Along the Spectrum

Existing law does not wait for full autonomy, and moving along the spectrum does not switch legal responsibility on or off. It changes where counsel should look for the legally significant choice and the evidence needed to explain it, as configurations, approval records, and logs gain importance.

> “Agentic AI may compress a workflow, but it does not compress the legal principles governing the outcome.

Recommend and draft. Regulation already attaches here. New York City’s law, for example, reaches tools that “substantially assist or replace discretionary decision making,” including screening.[23] Key questions include how a recommendation was produced and used by human decisionmakers, and whether agent-drafted content the employer adopts is accurate and consistent across documents.

Execute with approval. The approval point takes center stage, and the degree of regulation may depend on whether human involvement is meaningful, which can turn on whether the reviewer knows how to interpret the output, reviews it with other relevant information, and has authority to change the decision.[24] What information was presented, what could be changed, and what the record shows about the review all matter.

Execute autonomously. Focus moves upstream to the employer's deployment, permission, and monitoring choices, and laws keyed to consequential decisions, discussed later in this Report, apply most directly. Theories treating AI vendors that perform delegated hiring functions as the employer's agents also come into play (see Section III).

The key is calibrated control matched to proximity to a consequential employment decision. Agentic AI may compress a workflow, but it does not compress the legal principles governing the outcome. A simple but useful governance question: before the agent acts, who has the authority, and after it acts, who can explain why?

C. Common Workplace Applications of Agentic AI

Employers are using agentic AI most heavily for high-volume, repeatable work rather than one-off tasks. That makes sense: an agent is most useful when the work follows a recognizable pattern. In that kind of work, the agent can plan a sequence of steps and move between systems, such as the human resources information system (HRIS), the applicant tracking system (ATS), and payroll. It can then finish the job without a person stepping in at each stage.

The key change is from AI that *assists* to AI that *acts*. Generative AI drafts a job posting and waits for a recruiter to review and publish it, whereas an agent can publish the posting, screen the applicants, book interviews, and send a human the exceptions it cannot resolve. The underlying technology may be similar, but the authority the employer gives its agent is vastly greater.

1. Recruiting and Screening

Recruiting is the most mature use case. Agents now handle hiring workflows from start to finish, including sourcing candidates, screening applications against set criteria, contacting applicants, scheduling interviews, and following up for missing documents. Talent leaders expect this to grow. In Korn Ferry's *Talent Acquisition Trends 2026* survey of more than 1,600 talent leaders, 52% said they plan to add autonomous AI agents to their teams in 2026.[25]

Legally, the key change rests in the identity of the decision-maker. Earlier screening tools *recommended* which candidates a recruiter should advance or reject. An agent may simply screen candidates out on its own, so the rejected applicants are never seen by a person.

2. Scheduling and Task Allocation

Scheduling covers two very different kinds of work:

- **Administrative scheduling** means coordinating interviews, booking meetings, and managing calendars. It is generally low risk. If an agent books a conference room at the wrong time, it creates an inconvenience, not a legal claim.
- **Workforce scheduling** is different. When an agent assigns shifts, distributes overtime, or routes tasks among employees, it directly affects hours and pay. For example, if night shifts carry a pay differential, an agent that fills night shifts simultaneously decides which workers earn the higher rate. The same is true of overtime and of any premium pay assignment.

Tasks which look clerical nevertheless carry legal consequences. An agent that rebalances schedules in real time may quickly run into wage and hour obligations, and might run afoul of predictive scheduling laws that require advance notice of schedule changes, or premium pay when changes are made late. Additionally, it is not difficult to imagine a scenario where an agent might also distribute desirable assignments in ways that create disparities.

3. Performance Management

Performance management agents create and monitor productivity metrics, system activity, and even employee communications to draft or score evaluations. Employers currently deploy performance management agents with varying levels of autonomy, including (1) drafting performance reviews for a manager to finalize; (2) flagging employees whose

metrics fall below a set threshold; and (3) starting coaching or disciplinary processes based on metrics and performance reviews created by the agent.

Each step gives the agent more control over the outcome. At the third level, the agent is effectively starting an employment action.

These agents require constant access to employee activity data, so they often double as monitoring tools. A tool intended to make reviews more efficient can end up collecting data on how employees work, communicate, and spend their time. That raises privacy and data access questions, including what data the agent may collect, who can see it, how long it is kept, and whether employees were notified of it.

4. HR Service Delivery

In HR service delivery, agents now fully resolve employee requests rather than just answer questions. They handle leave, benefits, policy questions, and onboarding. For example, when an employee asks for leave, an agent can check eligibility, update the HRIS, notify payroll, and confirm the outcome to the employee, with no HR professional required in the process.

As a result, the agent may be the first to receive accommodation requests, leave requests, and internal complaints. Because employees are unlikely to label these requests precisely, an employee who tells the agent about a medical condition while asking about a schedule change may have triggered the employer's obligations, irrespective of whether the agent recognizes it.

> "Across all six applications, the agent's authority has grown, not just its capability.

5. Customer Service

Customer service is one of the most visible uses of agentic AI. Agents answer and resolve customer questions and pass complex or sensitive issues to human staff, and the pressure to adopt is strong. In a Gartner survey of 321 customer service and support leaders, 91% said executive leadership was pressing them to implement AI; nearly 80% plan to move at least some agents into new roles; and 84% plan to add new skills to the agent role and change their hiring profiles.[26]

The legal questions implicated are less about how the agent treats customers and more about what happens to the workforce around it. Potential employment effects include redesigned roles, reskilling programs, new performance measures for staff whose work is now supervising AI rather than handling customer contacts directly, and possible restructuring. Each brings its own considerations, including how jobs are classified, how performance standards are set for new roles, and the procedural requirements that apply to reductions in force.

6. Compliance Monitoring

Agents are increasingly used to monitor priorities such as timekeeping, safety, policy compliance, communications, and data movement. But unlike earlier monitoring tools that produced reports for later review, these agents act in real time. They flag, escalate, or restrict activity as it happens, creating a specific risk that the agent may detect protected activity or information. While scanning communications for policy violations, it may come across employee complaints about discrimination, coworkers discussing pay or working conditions, or reports of suspected fraud or safety violations.

When that happens, two questions follow:

- **Employer knowledge:** whether the employer is now considered to *know* about the protected activity, which can trigger employee protections or a duty to act and may affect a later retaliation claim.
- **NLRA Section 7:** whether monitoring, flagging, or restricting the activity interferes with employees' right to engage in concerted activity. *See* Section IV.F.1.

7. What These Applications Have in Common

Across all six applications, the agent's *authority* has grown, not just its capability. Screening tools could already rank candidates, and scheduling software could already build shift patterns. What is new is that employers increasingly let the agent act on its output without waiting for a person.

An employer's legal exposure thus depends less on what the agent can do than on the level of autonomy at specific decision points. An agent that drafts a review creates different risks than one that starts discipline, and an agent that recommends rejecting a candidate creates different risks from one that rejects the candidate itself.

D. Why Adoption Is Accelerating and What Makes It Durable

Agentic AI systems can complete multistep tasks with limited human direction, such as retrieving records, sending messages to employees or candidates, and updating information in connected systems. These capabilities offer efficiency gains and require employment counsel to consider how much autonomy the system will have. Whether agentic AI adoption is durable depends largely on data readiness, oversight design, and change management, and those same choices frame the legal issues addressed in Section IV.

Adoption is expanding, but different measures capture different developments. McKinsey's 2026 global survey reports that nearly nine in ten respondents' organizations use AI in at least one business function; among large enterprises, the share scaling AI agents rose from 27% to 40% over the preceding year.[27] It helps to distinguish four stages of change. **Deployment** means making a tool available. **Usage** means interacting with it. **Adoption** means integrating the tool into recurring work. **Transformation** means redesigning workflows, roles, or decision rights around it. Legal questions can arise at any stage, depending on the data processed, actions authorized, and effects on workers.

1. Three Pressures Driving Adoption

Three forces are pushing employers toward agentic AI: the prospect of faster work, pressure to reduce costs, and concern about falling behind competitors. Each rests on a different kind of evidence, and the evidence is often thinner than the enthusiasm. The discussion below separates what studies have measured from what surveys report and forecasts predict. That distinction matters for counsel, because the excitement and pressure behind a deployment often shape how quickly it moves and how carefully it is reviewed.

Efficiency and Productivity. The question for adoption, then, is not whether a system performs well in general but whether it reliably performs the specific work at issue under the conditions in which it will actually be used. A tool that excels at drafting, for example, may not be reliable when it makes decisions or operates connected systems. For counsel, that means asking for evidence of performance on the actual workflow, not general benchmarks or demonstrations.

Cost Pressure. Looking ahead, Gartner forecasts that agentic AI will resolve 80% of common customer-service issues without human intervention by 2029 and reduce operational costs by 30%.[28] The economic case depends on more than model or licensing charges. Integration, supervision, exceptions, correction, and rework can affect net savings. Measuring the cost of a successfully completed task, including human review, may therefore be more informative than counting automated steps. Counsel reviewing a business case can ask whether oversight, correction, and the cost of errors are included in the projected savings. Decisions about staffing and role changes raise additional issues addressed in Section IV.

Competitive Pressure. Competitive concerns may accelerate experimentation. In EY's 2025 technology-sector survey, 69% of technology leaders identified staying competitive as a reason for adopting agentic AI, and 92% planned to increase AI spending.[29] Those numbers explain why technology companies are moving quickly, but they say little about whether the same pace or deployment model suits every employer. Counsel can help separate commercial urgency from the time needed to assess a proposed use and any applicable obligations.

2. Foundations for Agentic Work

Data quality, completeness, provenance, and permitted use can affect both agentic AI performance and legal risk. Connecting an agent to personnel records also raises questions about which records it needs, who may receive its outputs, and whether it can modify or disclose information. An agent that adjusts shift schedules, for instance, may need availability data but not disciplinary history. And because each step builds on the last, a single error can carry through the rest of the workflow.

At the system level, one approach is to match an agent's permissions to its task, treating access to read information separately from authority to change records, send communications, or initiate consequential actions. Organizations should also plan how to pause a workflow, route an exception to a responsible person, and correct or reverse an action where feasible. NIST's voluntary AI Risk Management Framework provides a useful reference for human oversight, monitoring, incident response, and recovery. It does not prescribe a single deployment model or establish compliance with employment law.[30]

Organizational fit matters too. Gallup found frequent AI use among 88% of employees who strongly agreed that AI integrated well with existing systems and processes, compared with 55% of those who did not strongly agree.[31] Gallup reported a similar pattern for manager support, with frequent use at 78% among employees who strongly agreed that their manager actively supports AI use, compared with 44% of those who did not.[32] The data show correlation rather than causation, and they concern AI use generally rather than agents. Still, they suggest that making a tool available is only the start.

3. Adoption Choices that Shape Legal Risk

Workflow selection and the allocation of human responsibility are particularly relevant to employment counsel. What follows is not a checklist; how much each point matters will depend on the system, its use, the workforce, and the governing law.

Workflow Selection. A workflow's legal significance depends on what it does and how its outputs are used, not on how it is labeled. Calling a tool "administrative" does not take it outside the law, and calling it "agentic" does not by itself bring it within any particular rule. An interview scheduling tool that only coordinates calendars, for example, sits further from a hiring decision than one whose output affects which candidates move forward in the process. Relevant considerations include proximity to employment decisions; the sensitivity and sources of data; the people affected; the jurisdictions involved; and the consequences and reversibility of error.

These considerations may support a pilot with guardrails, a narrower grant of authority, or a different point of human intervention. They can also inform review of vendor responsibilities and the information available to evaluate and explain the system's role. Some laws turn on whether automated output substantially assists a decision or whether a decision is solely automated, rather than whether the employer calls the technology an agent. New York City's automated employment decision tool framework and General Data Protection Regulation (GDPR) Article 22 illustrate these different approaches. Their definitions, exceptions, and territorial reach require separate analysis.[33]

Human Involvement. Beyond the system-level controls described in Section B, human involvement can occur before an action, at defined checkpoints, or when an exception arises. Its practical value depends on the reviewer's information, competence, time, and authority to question the output, change the result, or stop the process.

For a particular workflow, counsel and operational teams should decide which actions warrant review before execution, which conditions trigger escalation, and who can intervene when the agent encounters missing information or unexpected instructions. An agent that drafts attendance warnings, for example, might queue them for a manager's

approval rather than send them directly. Records of the inputs, material actions, exceptions, and human decisions may help explain what occurred, subject to applicable privacy, retention, and preservation requirements. The weight given to human review will depend on the governing law and the facts of the particular process.

AI Literacy and Use. Agentic AI changes how work gets done, who does it, and who answers for the result. Organizations that sustain use treat adoption accordingly, as a change-management effort rather than a software installation. Without that effort, tools are often deployed but little-used, or used in ways no one designed or reviewed. That second outcome carries legal as well as operational risk: the permission limits and review checkpoints described above work only if the people around the agent understand and follow them.

Effective change generally has to come from two directions. From the top, senior leadership needs to want the change and fund it, and leaders who use the tools themselves and explain the purpose of the change can help set expectations. From the bottom up, training programs, peer-to-peer learning, and champion networks help employees build the skill and confidence to use the tools well. Recognizing effective use can help new practices take hold.

Training works best when it focuses on the work employees actually perform, such as recognizing limitations, evaluating outputs, using approved tools, and knowing when and how to escalate. Clear ownership matters as well: someone should be accountable for each workflow's performance, exceptions, and changes over time.

Evaluation should distinguish activity from outcomes. Alongside usage, an employer can examine task completion, output quality, error severity, exceptions, review time, and correction costs against a relevant baseline. Testing ordinary cases and foreseeable edge cases (i.e., unusual cases where a tool may fail to work as expected) may reveal whether apparent time savings survive human review and whether a change in the model, data, or permissions changes the result.

Communication is part of the change effort. Employees will want to know how AI affects their jobs and where they can raise concerns, and gaps in information tend to fill with speculation. Communications about a pilot can distinguish its current scope from possible future changes without promising outcomes not yet determined. Workforce restructuring, employee monitoring, and potential consultation or bargaining obligations are addressed in Section III.

4. Expansion, Retrenchment, and Delegated Authority

Most forecasts expect agents to spread quickly.[34] However, it is unclear how much authority employers will give agents, what productivity they will realize, or how many jobs will change. For employment counsel, delegated authority is the development to watch, and the path is unlikely to be a straight line.

Expansion and retrenchment may occur together.[35] A useful assessment therefore asks what the agent is authorized to do, what evidence supports that use, where people can intervene, and what happens when the process fails. Those same choices determine whether adoption lasts, and they provide the factual setting for the discrimination, privacy, labor, accountability, and regulatory questions examined in Section IV.

II. The Current Regulatory Environment for Workplace AI

Before analyzing the legal risks posed by use of autonomous AI agents in employment functions, it is necessary to look at the regulatory backdrop: the laws that have been passed globally that address how an employer may use AI systems.

A. The U.S. Federal Landscape

Although Congress has not enacted a comprehensive federal regulatory framework governing workplace AI, federal labor and employment agencies have increasingly addressed AI-related issues through guidance documents, enforcement initiatives, strategic plans, public statements, and administrative actions. The Equal Employment Opportunity Commission (EEOC), National Labor Relations Board (NLRB), and U.S. Department of Labor (DOL) have each taken distinct approaches reflecting their statutory mandates. For employers deploying agentic AI, the principal federal message has been consistent: artificial intelligence may be a new technology, but it does not create exemptions from longstanding obligations concerning discrimination, collective bargaining, employee rights, wages, workplace safety, privacy, and transparency.

1. EEOC: Applying Civil Rights Laws to AI-Based Employment Decisions

Throughout changes in leadership and presidential administrations, the EEOC has consistently taken the position that existing federal antidiscrimination laws apply fully to AI-driven employment decisions. While a number of sub-regulatory documents issued during the prior administration concerning the discriminatory use of AI have been withdrawn, the agency's overarching policy guidance[36] remains online as of this writing.

In its guidance, the EEOC's focus has centered on the use of AI in recruiting, screening and hiring, as well as workplace surveillance, promotion, compensation, performance management, and termination decisions. The agency has previously emphasized that reliance on an algorithm does not relieve an employer of responsibility under federal civil rights statutes, much in the way the use of a discriminatory physical or pen-and-paper traditional test would not shield the employer from a claim of unlawful discrimination. Under the current administration, the EEOC generally does not appear to have made algorithmic discrimination a priority. Whether that will continue to be true for the remainder of the administration remains unclear.

Much of the agency's deemphasis is likely due to the fact that many discrimination claims arising from the use of artificial intelligence are based on a "disparate impact" theory of liability. Put simply, disparate impact cases may arise where the use of an AI tool with an objective that is facially neutral results in a disproportionate negative impact on individuals of a certain protected class. This does not mean use of the tool is per se illegal, as it may be justified as the least discriminatory method justified by business necessity. Following on an executive order[37] issued near the start of the administration, the EEOC has dramatically de-emphasized the use of disparate impact theory and closed almost all

pending charges based solely on allegations of disparate impact.[38] In its most recent National Enforcement Plan—a document laying out the agency's substantive priorities and areas of focus—AI and algorithmic discrimination is not mentioned at all. And, as of this writing, the only lawsuit the EEOC has brought in federal court involving the use of an AI tool alleged that the tool was *intentionally* discriminatory and purposefully programmed to screen out older applicants.

Notwithstanding the change in the federal enforcement landscape, the underlying statutory framework remains unchanged. Title VII, the Americans with Disabilities Act, the Age Discrimination in Employment Act, the Genetic Information Nondiscrimination Act, and other federal anti-discrimination laws continue to apply to AI-assisted employment decisions. As a result, employers using agentic AI remain exposed to litigation risk even if federal enforcement priorities evolve. Private plaintiffs, state regulators, and courts may likewise continue to challenge biased AI systems under existing employment statutes.

2. NLRB: AI Through the Lens of Collective Bargaining and Employee Rights

The National Labor Relations Board (NLRB) has not yet established jurisprudence specific to agentic AI. When it does face questions about whether and how agentic AI impacts employee rights, collective bargaining, and employers' obligations under the National Labor Relations Act (NLRA), it is expected that the decisions will follow traditional labor law principles governing automation and technological change. Future cases will also require contemplation of whether labor law doctrine developed by evaluating human actors and decision-makers should apply when AI has replaced those human actors, in whole or in part.

Electronic Monitoring and Surveillance. Modern agentic systems can monitor productivity, track employee movements, analyze communications, and evaluate performance. Agents that monitor, score, or review employee communications can operate continuously and at scale creating uncertainty about whether protected activity is being captured or could reasonably appear to be. Unlike passive tools, an agent may also act on what it observes (flagging, escalating, or adjusting work, or referring employee conduct for investigation or discipline) without contemporaneous human review. Unions may argue that monitoring, even by a non-human system, chills employees' exercise of their rights to organize and engage in protected concerted activity and is violative of the Act. Regardless of whether this turns out to be true or not, employers should consider whether the appropriate safeguards are in place in connection with any monitoring so the organization captures the value of the information collected without infringing on employee rights.

Bargaining Obligations. When employers introduce technologies that affect terms and conditions of employment, planning should include whether decision and/or effects bargaining will be required. Agentic AI may affect job duties, staffing levels, scheduling, workplace policies, performance evaluation criteria, discipline systems, workload allocation and more. Depending on the impacts, existing contract language and related considerations, employers may be required to bargain over the decision to implement agentic AI, the effects of implementation, or both. If bargaining is contemplated, employers can also plan thoughtfully around the form and timing of any such negotiations.

AI and Workforce Displacement. As agentic AI systems become capable of performing increasingly sophisticated white-collar and administrative functions, concerns regarding job displacement have intensified. Unions are using this fear to recruit white-collar workers to consider union representation in increasingly large numbers. It has also led to an increased focus on AI and displacement concerns at the bargaining table. Labor organizations have sought contractual protections limiting AI deployment, and where jobs are at risk, retraining opportunities, severance benefits, and advance notice requirements. Whether or not a bargaining obligation exists and how contract language will be interpreted will mostly likely be analyzed under traditional labor law principles governing automation and technological change. Section IV.F discusses how these principles may apply to agentic AI.

3. DOL: Prioritizing AI to Increase Productivity and Economic Growth

The DOL under the second Trump administration has taken a substantially different approach to artificial intelligence than the Biden administration, emphasizing innovation, economic competitiveness, and workforce development rather than AI-specific workplace regulation. Consistent with President Trump's executive order on removing barriers to American leadership[39] in AI and the administration's subsequent AI Action Plan,[40] the Department has characterized AI primarily as

an opportunity to increase productivity, create new jobs, strengthen national competitiveness, and support economic growth. Public statements from DOL leadership have stressed that American workers should be equipped to harness AI rather than be protected from AI through new regulatory requirements.[41] This emphasis is expected to continue under recently confirmed Secretary of Labor Keith Sonderling, who in his prior position as deputy secretary of labor, and previously as a commissioner of the EEOC, took a keen interest in AI issues and how to best maximize their utility in the workplace.

A central pillar of the DOL's AI agenda is workforce preparation. Under the administration's AI Action Plan, the Department has been assigned responsibility for initiatives designed to "Empower American Workers in the Age of AI" and "Train a Skilled Workforce for AI Infrastructure."[42] To accomplish these goals, the Department has promoted AI literacy, encouraged the incorporation of AI-related training into workforce development programs, expanded apprenticeship opportunities, and supported the development of career and technical education programs focused on AI-related skills. The Department has also encouraged state and local workforce agencies to use existing federal workforce funds to support AI education and digital-skills development, reflecting a view that the principal labor-market challenge posed by AI is ensuring workers have the skills necessary to adapt to technological change.

B. U.S. State and Local Regulation

1. Overview of State and Local Laws

In December 2021, New York City enacted a first-in-the-nation law regulating the way employers may use AI tools in their employment decisionmaking. Each year since, an increasing number of U.S. jurisdictions have introduced similar legislation. In December 2025, President Trump signed an executive order entitled “Ensuring a National Policy Framework for Artificial Intelligence,” which attempts to limit states’ ability to regulate the use of AI by establishing an AI Litigation Task Force within the Department of Justice that is authorized to challenge state AI laws on the grounds that they are unconstitutional burdens on commerce, are preempted by federal regulations, or are otherwise unlawful.[43] Nonetheless, nearly half of U.S. states introduced AI-related legislation during the 2026 legislative session, and it is anticipated that state legislatures will continue to take up the issue of AI regulation in 2027.

Following the enactment of New York City’s law, a few other jurisdictions have passed laws regulating AI in employment. The tables below provide a brief overview of these laws. A full description of each is included in Appendix A.

Jurisdiction and Title	Synopsis
California S.B. 947 (2026)	Prohibits an employer from using automated decision systems (ADS) or relying solely on ADS in making employment decisions; requires an employer to provide post-use notice of ADS to employees. Effective July 1, 2027.
California A.B. 1883 (2026)	Prohibits employers from using a workplace surveillance tool that uses artificial intelligence to collect neural data or recognize an individual’s emotional state. Effective January 1, 2027.
California S.B. 951 (2026)	Requires employers to provide a specific technology-related Cal-WARN notice before separating workers if the separation is due to displacement by an AI system or other automated technology. Effective January 1, 2027.
California Civil Code	Provides that in an action against a defendant who developed, modified, or used AI that is alleged to have caused harm to the plaintiff, the defendant cannot assert that the AI autonomously caused the harm to the plaintiff. Currently in effect.

Jurisdiction and Title	Synopsis
California Fair Employment and Housing Act (FEHA) Regulations	Makes it unlawful for an employer to use an ADS or selection criteria (including a qualification standard, employment test, or proxy) that discriminates against a job applicant or an employee or a class of applicants or employees on a basis protected under the FEHA, subject to any available defenses. Currently in effect.
California Consumer Privacy Act Regulations	Requires covered entities to conduct risk assessments if using automated decisionmaking technology; creates data access and opt-out rights for consumers related to a business's use of automated decisionmaking technology. Currently in effect.
Colorado Automated Decision-Making Technology Act	Regulates the use of automated decision-making technology (ADMT) to materially influence a consequential decision, which includes access to, eligibility for, selection for, or compensation for employment. Effective January 1, 2027.
Connecticut S.B. 5 (2026)	• Requires an employer to notify applicants and employees if an AI system will be used in making employment-related decisions; prohibits use of an AI system to make discriminatory employment-related decisions. Effective October 1, 2027. • Requires an employer conducting a WARN-covered layoff to disclose to the state labor department whether the layoffs are related to the employer's use of AI or another technological change. Effective October 1, 2026. • Provides that an employer's use of an automated employment decision tool is not a defense to a discrimination complaint. Effective October 1, 2026.
Connecticut S.B. 472 (2026)	Requires employers to notify employees of electronic surveillance and which prohibited activities may require surveillance without prior notice. Currently in effect.
Illinois Human Rights Act	Prohibits an employer from using AI functionality that has the effect of subjecting employees and applicants to discrimination on the basis of a protected class in connection with the following facets of the employment lifecycle: recruitment, hiring, promotion, renewal of employment, selection for training or apprenticeship, discharge, discipline, tenure, or the terms, privileges, or conditions of employment. Currently in effect.
Illinois Artificial Intelligence Video Interview Act	Provides that an employer that asks applicants to record video interviews and uses an artificial intelligence analysis of applicant-submitted videos must notify each applicant in writing before the interview. Currently in effect.
Maine L.D. 61 (2026)	Requires employers that use employee surveillance to provide notice to employees and candidates for employment. Currently in effect.
New Jersey Law Against Discrimination Regulations	Provides that the use of automated decisionmaking technology to make employment decisions, including but not limited to decisions related to advertising, recruiting, screening, interviewing, hiring, and compensation, may give rise to a disparate impact on members of a protected class. Currently in effect.
New York State Worker Adjustment and Retraining Notification (WARN) Act	Requires WARN Act notices to disclose whether the layoffs are related to the employer's use of AI.
New York City Local Law 144 of 2021	Requires an employer that uses an automated employment decision tool to screen a candidate for an employment decision to notify each such candidate regarding use of the AI tool; requires an independent bias audit on any such tool in use. Currently in effect.

The following table provides the statutory and regulatory definitions for the technologies covered under each of the laws and regulations discussed above. The definitions provide a foundation for analyzing how agentic AI, and specific AI agents and their functions within a workflow, fit within the current regulatory environment.

Jurisdiction	Law/ Regulation	Term/Covered Technology	Definition
California	Labor Code §§ 1520 *et seq.*	Automated decision system	Any computational process derived from machine learning, statistical modeling, data analytics, or artificial intelligence that issues simplified output, including a score, classification, or recommendation, that is used to assist or replace human discretionary decisionmaking and materially impacts natural persons.
California	Labor Code §§ 1400.5 (CalWARN), 1520 *et seq.* (No Robo Bosses Act)	Artificial intelligence	An engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments. An ADS does not include a spam email filter, firewall, antivirus software, identity and access management tools, calculator, database, dataset, or other compilation of data.
California	Labor Code §§ 1580 *et seq.*	Workplace surveillance tool	Any system, application, instrument, or device that collects or facilitates the collection of employee data, activities, communications, actions, biometrics, or behaviors by means other than direct observation by a person, including but not limited to video or audio surveillance, continuous incremental time-tracking tools, geolocation, electromagnetic tracking, photoelectronic tracking, or that utilizes a photo-optical system or other means.
California	FEHA Regulations	Automated decision system	A computational process that makes a decision or facilitates human decisionmaking regarding an employment benefit.
California	CCPA Regulations	Automated decisionmaking technology	Any technology that processes personal information and uses computation to replace or to substantially replace human decisionmaking.
Colorado	Colorado AI Act	Automated decisionmaking technology	Technology that processes personal data and uses computation to generate output, including predictions, recommendations, classifications, rankings, scores, or other information that is used to make, guide, or assist a judgment, decision, or determination concerning an individual.
Connecticut	SB 5 (not yet codified)	Automated employment decision tool	Any technology that processes personal data and uses computation to generate any output, including, but not limited to, any prediction, recommendation, classification, ranking, score or other information, that is a substantial factor used to make or materially influence an employment-related decision.

Jurisdiction	Law/ Regulation	Term/Covered Technology	Definition
Connecticut	SB 5 (not yet codified)	Artificial intelligence	Any machine-based system that, for any explicit or implicit objective, infers from the inputs such system receives how to generate outputs, including, but not limited to, content, decisions, predictions or recommendations, that can influence physical or virtual environments.
Illinois	Illinois Human Rights Act	Artificial intelligence	A machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments, and includes generative artificial intelligence.
Illinois	Artificial Intelligence Video Interview Act	Artificial intelligence	No definition provided
New Jersey	NJLAD Regulations	Automated decisionmaking tool	Any technological tool that aims to automate all or part of the human decisionmaking process.
New York City	N.Y.C. Admin. Code §§ 20-870 *et seq*.	Automated employment decisionmaking tool	Any computational process, derived from machine learning, statistical modeling, data analytics, or artificial intelligence, that issues simplified output, including a score, classification, or recommendation, that is used to substantially assist or replace discretionary decision making for making employment decisions that impact natural persons.

2. Extraterritorial Application of U.S. AI Laws

As a general rule, there is a presumption against extraterritorial applicability of a state's laws absent a clear intent to the contrary.[44] The Commerce Clause "precludes the application of a state statute to commerce that takes place wholly outside of the state's borders, whether or not the commerce has effects within the state."[45] A state may not project its law into another state to control commerce or transactions occurring there, but it may apply its law to multistate conduct with substantial in-state connections or effects if the regulation is not discriminatory, does not directly control out-of-state commerce, and does not impose burdens on interstate commerce that are clearly excessive in relation to the putative local benefits.[46]

Using the example of a California employer utilizing an AI resume screening tool to find candidates for a job vacancy that was widely published online on websites with a national reach, a state may likely regulate employers' use of such tools when the regulation is evenhanded, tied to hiring with a meaningful in-state nexus, and aimed at local interests such as fair employment and applicant protection. Under that model, any effect on nationwide recruiting is usually treated as incidental and judged under the burdens on commerce vs. local benefits paradigm, not invalidated automatically.[47] However, the law becomes constitutionally vulnerable if its practical effect is to force employers or AI vendors to conduct out-of-state screening according to one state's terms for transactions or jobs wholly outside the state.

C. Global Framework

The European Union, Canada, the United Kingdom, China, South Korea, and Brazil have workplace AI laws addressing human oversight, explanation, contestability, and territorial reach in employment-related AI. None of these jurisdictions has enacted binding legislation specific to agentic AI. China has, however, issued a dedicated policy document on AI agents, and regulators elsewhere have published non-binding analyses. The six jurisdictions follow different models:

The EU and South Korea have enacted comprehensive, risk-based AI statutes that differ considerably in scope and intensity.

China regulates through layered, technology- and service-specific administrative rules alongside its data protection law.

The United Kingdom, Canada, and Brazil have no AI-specific statute in force and rely primarily on data protection law, sectoral or provincial rules, and regulatory guidance. Canada and Brazil have bills pending.

Obligations depend on the system's function, risk, classification, processing of personal data, deployment context, and effects on individuals. A more detailed description of the above-mentioned laws is included in Appendix A.

1. Human Oversight and Explainability vs. Autonomous Operations

The central issue is not whether autonomous operation is prohibited, but whether the agentic process remains attributable to identifiable persons and organizations as operational autonomy increases. For agentic AI, the key questions are where in a multi-step workflow a legally relevant decision occurs and which oversight, explanation, or review mechanism attaches at that point. The jurisdictions answer these questions differently. Some prohibit solely automated significant decisions in principle (the EU); some permit them subject to safeguards (the United Kingdom); and others grant individuals rights to an explanation, review, or refusal on request (China, South Korea, Québec, and Brazil).

European Union: AI Act. The EU framework distinguishes system-level oversight from individual rights. Article 14 of the EU AI Act requires covered high-risk systems to be designed so that natural persons can effectively oversee their operation. The system must enable the persons assigned to oversight, as appropriate, to understand its capacities and limitations. Designated human supervisors remain alert to automation bias, interpret, disregard, or override outputs, and interrupt the system. Article 26 requires deployers to assign oversight to persons possessing the necessary competence, training, authority, and support, and requires employers to inform workers' representatives and affected workers before putting a high-risk system into service at the workplace. Article 86 grants individuals affected by significant decisions based on the output of a high-risk system (under Annex III of the EU AI Act) a right to an explanation of the system's role, unless Union law already provides one. These provisions do not require a person to make every AI-assisted employment decision, and they apply to Annex III systems only from December 2, 2027.

European Union: GDPR Article 22. Article 22 of the GDPR addresses a different risk. It prohibits, in principle, decisions based solely on automated processing that produce legal or similarly significant effects, subject to exceptions for contractual necessity, authorization by Member State law, and explicit consent, which is rarely a reliable basis in the employment context due to the perceived imbalance of power. Where an exception under Article 22(2)(a) or (c) applies, safeguards include the right to obtain human intervention, express a viewpoint, and contest the decision. The Court of Justice of the European Union (CJEU) has held that an automated score can itself constitute a decision where a third party draws strongly on it, and that the right of access extends to the procedure and principles actually applied, subject to the protection of trade secrets. Under guidance endorsed by the European Data Protection Board (EDPB), human involvement takes a decision outside Article 22 only where the reviewer has the authority and competence to change it and considers all relevant information. Where a human merely ratifies an agent's ranking or recommendation, the agent's output may itself be the relevant decision. Operational oversight under the EU AI Act and individual intervention under the GDPR are therefore related but legally distinct mechanisms.

European Union: Platform Work Directive. For persons performing platform work, the Platform Work Directive, once transposed (by December 2, 2026), will require human monitoring, human review of significant decisions, and that account suspensions or terminations be decided by a human being.

> "For multinational employers, the territorial question is not where an agent is hosted but where its effects land.

United Kingdom. In the United Kingdom, the Data (Use and Access) Act 2025 amended the existing automated decision-making provisions of UK GDPR and inserted new Articles 22A–22D into UK GDPR, with effect from February 5, 2026. A decision is solely automated where there is no meaningful human involvement. Significant solely automated decisions are generally permitted subject to Article 22C safeguards: information, an opportunity to make representations, human intervention, and contestation. Article 22B restricts such decisions based entirely or partly on the processing of special-category data unless a condition applies, which is unlikely to be the case in the employment context. The Information Commissioner's Office (ICO) 2026 report on agentic AI sets out its early views rather than guidance; the ICO has announced updated guidance on automated decision-making.

China. China's Personal Information Protection Law (PIPL) requires personal information handlers using automated decision-making to ensure transparency and fair and impartial results. When a decision made solely through automated decision-making significantly affects an individual's rights and interests, the individual may request an explanation and refuse that method of decision-making. Handlers must also conduct a prior personal information protection impact assessment. Rules on algorithmic recommendation services add protections for workers subject to platform work-scheduling algorithms. China's 2026 policy guidance on AI agents, which is not legally binding, calls for users to retain the right to be informed of, and to make the final decision on, an agent's autonomous decisions.

South Korea. South Korea's AI Basic Act requires operators of high-impact AI to notify users in advance that a product or service is based on such AI and to implement measures addressing risk management, explanation of outcomes and principal criteria, user protection, human oversight, and documentation. Documentation with explanations is required to the extent technically feasible. Whether an employer that uses a third-party system solely for internal HR purposes is an operator subject to these duties has not yet been settled. Article 35 frames the fundamental-rights impact assessment as a best-efforts obligation. The AI Basic Act's operator categories do not map directly onto the EU concepts of provider and deployer. Individual decision rights instead derive from the Personal Information Protection Act (PIPA): data subjects may, subject to exceptions, refuse fully automated decisions that significantly affect their rights or obligations and request an explanation, and controllers must publish the criteria and procedures used.

Canada. Canada presents a sectoral and provincial model. The federal Directive on Automated Decision-Making requires covered federal institutions to publish algorithmic impact assessments and implement impact-scaled explanation, human-involvement, and recourse measures. Québec requires an enterprise making a decision based exclusively on automated processing to notify the affected person and, on request, disclose the personal information, reasons, and principal factors involved. The person must be allowed to submit observations to a staff member able to review the decision. In Ontario, since January 1, 2026, employers with at least 25 employees must disclose AI use to screen, assess, or select applicants in publicly advertised job postings. The rule requires disclosure, not human review or an explanation of the system. For agentic workflows, Québec's "exclusively automated" test raises the same question as GDPR Article 22: whether a staff member who simply approves an agent's ranking or recommendation takes the decision outside the rule, and Ontario's disclosure duty may apply wherever an agent screens or ranks applicants for an advertised position.

Brazil. Article 20 of Brazil's General Data Protection Law (LGPD) permits a data subject to request review of decisions made solely through automated processing that affect the subject's interests and to obtain information about the criteria and procedures used, subject to commercial and industrial secrecy. The enacted text does not expressly require review by a natural person. PL 2338/2023 would classify AI used in recruitment, evaluation, promotion, and termination as high-risk

and establish broader AI-specific rights, but remains subject to amendment. For agentic workflows, the review right may attach where an agent's output is adopted without independent human assessment. While the LGPD does not require the review to be performed by a person, it remains legally unsettled whether, and in what circumstances, a second automated check would satisfy this requirement. If enacted in its current form, PL 2338/2023 would likely treat agents used for those employment functions as high-risk.

These regimes do not use oversight, intervention, review, explanation, and contestability interchangeably, so each mechanism should be matched to its legal trigger, responsible actor, and decision stage. Explanation and contestation of significant decisions are common to all six jurisdictions and regulators generally treat human review as meaningful where the reviewer has the authority, competence, and information needed to change the outcome. Employers may therefore consider identifying the points at which an agent's output becomes determinative and documenting the oversight applied there.

2. Extraterritorial Reach of Non-U.S. Laws

For multinational employers, the territorial question is not where an agent is hosted but where its effects land. Each regime uses its own connecting factors, such as establishment, the targeting or monitoring of individuals, the use of a system's output, or effects on the local market, and employment protections attach separately to the place of work. An agent hosted in the United States that screens applicants for a role in Germany, for example, may bring the EU AI Act (because its output is used in the Union), the GDPR, and German works council rights into play at once.

Jurisdiction	Law	What brings a foreign employer within reach	Local representative
European Union	EU AI Act	Providers placing systems on the Union market; deployers established or located in the Union; non-EU providers and deployers where the system's output is used in the Union	Non-EU providers of high-risk systems must appoint an authorized representative in the Union
European Union	GDPR	Processing in the context of an EU establishment; specified offering or monitoring activities directed at individuals in the Union	Non-EU controllers/processors caught by Art. 3(2) must appoint a representative in the Union (exceptions for occasional/low-risk processing and public authorities)
United Kingdom	UK GDPR (Art. 3)	Establishment and targeting tests comparable to the GDPR	Non-UK controllers/processors caught by Art. 3(2) must appoint a representative in the UK, subject to the same occasional/low-risk and public-authority exceptions as under the EU GDPR
China	PIPL	Specified overseas processing of personal information concerning persons in China, including processing to provide products or services or to analyze or evaluate behavior	Overseas handlers must designate a representative in China
South Korea	AI Basic Act	Conduct abroad that affects the Korean market or users within the Republic of Korea	Domestic agent required for certain foreign operators (see note below)
Canada	PIPEDA and provincial statutes	Real and substantial connection to Canada; employee-data coverage is limited (see note below)	No local representative requirement under PIPEDA, Ontario statutes, or Quebec's Law 25

Jurisdiction	Law	What brings a foreign employer within reach	Local representative
Brazil	LGPD	Processing occurs in Brazil; processing targets the provision of goods or services to, or concerns data of, individuals located in Brazil; or the data was collected in Brazil	No local representative requirement

Note on Canada and South Korea. For employee data, PIPEDA reaches only federal works, undertakings, or businesses; other employees are covered, if at all, by provincial statutes such as those of Québec, Alberta, and British Columbia. Under South Korea's AI Basic Act, foreign operators without a Korean place of business must designate a domestic agent if they exceed revenue or user thresholds.

Employment law layer. Data and AI rules are only one layer. Under EU conflict-of-laws rules, employees may not be deprived of the mandatory protections of the law of the country where they habitually work, and local consultation and co-determination rights, such as German works council rights over technical monitoring systems, attach to the local establishment regardless of where the system is hosted.

For multinational employers, centralized procurement, hosting, or model ownership therefore does not alone determine territorial application. Employers may address overlapping requirements through a group-wide baseline, jurisdiction-specific configurations, or a combination of both.

III. Reconciling the Current Regulatory Environment with the Advances in Autonomous AI Systems

A. Mismatch: AI Employment Laws and Agentic Systems

The explainability critique of AI is familiar: complex machine-learning models operate as "black boxes," and even their developers may be unable to say why one applicant scored higher than another. Regulators often respond by demanding transparency, documentation, human review, or disclosure of decision factors. The familiar response is that human decision-making is hardly more transparent. Managers rely on intuition and pattern recognition they cannot fully articulate, their explanations are often post hoc, and the law tolerates considerable opacity in human decisions.

Defensibility. The more serious problem is not explainability but defensibility. Discrimination law has never required perfect explanations from human decision-makers; instead, it tests whether outcomes can be justified. That is where the conflict becomes concrete. Disparate impact doctrine rests on a causal, hypothesis-based model of selection, while machine learning rests on correlation-based prediction that need not be tied to causation.

Content, construct, and criterion validity. Once a plaintiff shows that a selection procedure disproportionately excludes a protected group, the employer must show that the practice is job-related and consistent with business necessity. The Uniform Guidelines on Employee Selection Procedures,[48] on which courts have relied for decades, assume that a procedure can be studied, understood, and defended through validation. Content validity ties the procedure to actual job tasks; construct validity ties it to a trait important to performance; and even criterion validity, which rests on a statistical relationship with job performance, traditionally operates against a backdrop of identifiable job characteristics. Each assumes the employer can say what the procedure measures and why it matters for the job.

> "The more serious problem is not explainability but defensibility.

Application to machine learning. Machine learning strains that assumption. A model may weigh thousands of variables and their interactions, identifying patterns that improve predictive accuracy without any underlying theory of job success. Suppose a hiring model finds that a combination of linguistic patterns, timing of career moves, educational attributes, and interview characteristics predicts sales performance, and its use creates a disparate impact. The employer may be able to show that the model predicts performance. What it may be unable to show is why the predictors matter. The model may know that something predicts success without knowing what job-relevant characteristic it is measuring, and "the model works" may not, standing alone, carry the employer's burden.

Criterion justification. The difficulty is not correlation itself; criterion validation has always relied on correlations. It is that business necessity calls for a justification of why a factor should be allowed to influence employment decisions despite its discriminatory effect. Employers have traditionally met that burden by showing that the factor measures something important to the job, such as judgment, conscientiousness, or technical competence. An advanced model's only answer may be that the data show the factor predicts success. Under current frameworks, that may not be enough.

Agentic systems sharpen the problem. An agent that adjusts its own criteria as it works may leave the employer unable to identify which factors drove a given outcome, let alone why, as the "connectedness" hypothetical in Section IV.C illustrates. And even where an employer meets its burden, a plaintiff may still prevail by showing a less-discriminatory alternative, an argument that may be easier to make when a model can be reconfigured.

A new question. AI presents a question that the existing legal framework was never designed to answer: if a system demonstrably predicts successful job performance better than human decision-makers, but cannot explain the causal basis for its predictions, can it ever fully satisfy a business-necessity defense? The Uniform Guidelines largely assume that the answer will come from validation grounded in identifiable job-related characteristics. Modern machine learning increasingly suggests that highly effective prediction may exist without any corresponding causal explanation.

The core tension. Anti-discrimination law is built around decision justification through human explanation. Advanced AI is built around decisional justification through computerized prediction, centering status-quo accuracy over required fairness. As AI grows more sophisticated and more decision-making authority shifts from humans to AI, these considerations become increasingly difficult to reconcile. With the advent of agentic AI, the ceding of autonomy is close to complete. Agentic systems may be able to conduct an entire employment process without human supervision or interdiction. At that point, the AI is the employer's agent for all practical purposes, and the core tension becomes even harder to reconcile, for two reasons.

Agentic AI further separates outcomes from identifiable selection criteria. Traditional disparate impact analysis generally assumes that we can identify the challenged practice. A cognitive test, a degree requirement, a personality assessment, a physical-strength standard, or even an AI-built scoring model can be isolated and validated. We can ask whether a particular criterion causes adverse impact, and whether it is job-related and consistent with business necessity. Agentic systems increasingly do not operate through a fixed and stable set of criteria. They dynamically gather information, formulate plans, decide what data is relevant, choose among alternatives, invoke tools, revise conclusions, and adapt their own workflows in pursuit of an objective. The question "which factor caused the disparate impact?" becomes far more difficult when the system itself is continuously deciding which factors to consider and how much weight to place on them. The legal framework assumes a relatively stable selection procedure. Agentic systems can function like an evolving decision rubric instead of a set company-wide policy.

Agentic AI widens the gap between prediction and justification. Even with conventional machine learning, an employer may struggle to explain why a particular model predicts success. With agentic AI, the difficulty increases because the agent may not merely be predicting an outcome, it may be exercising judgment. For instance, an agent may review applications, investigate candidates, search for additional information, determine which qualifications are most important for a specific opening, deprioritize some criteria, emphasize others, and ultimately recommend disposition decisions. At that point, the AI is no longer merely finding correlations within a predefined model for humans to use as part of their decisional matrix. Instead, it is making a series of intermediate judgments that collectively produce the outcome. If those outcomes generate adverse impact, employers' burden under traditional disparate-impact principles becomes extraordinarily difficult. It is not enough to say that the final recommendation predicts success. Employers would need to explain why the various considerations the agent deemed relevant are job-related and consistent with business necessity. The more autonomous the system becomes, in other words, the harder that explanation becomes, and full autonomy is the agentic AI goal.

B. The "No Robo Bosses" Era: Employment Law in the Age of Agentic AI

1. From AI Assistant to AI Agent

A few years ago, workplace artificial intelligence was most visible at the front end of the employment relationship. Employers used automated tools to search resumes, identify keywords, and sort applicants into manageable pools. Those systems could be consequential, but their role was generally discrete: they performed a defined task and returned an output to a human decision-maker. Agentic AI changes that model. Instead of waiting for a person to direct each step, an AI agent can pursue an assigned objective across multiple stages of a workflow. The relevant legal question is therefore no longer only whether an algorithm screened a resume fairly. It is whether an automated system has become an active participant in a chain of employment decisions.

That shift creates an accountability problem. When an AI agent influences a discriminatory hiring result, an erroneous disciplinary recommendation, or a reduction in force, who bears responsibility: the employer, the software developer, the vendor that configured the system, or the individual manager who accepted the output? The emerging answer from legislatures and regulators is increasingly clear. An employer may use AI to assist its decisions, but it cannot delegate legal accountability to the machine.

The newest laws do not follow a single model. Some require audits. Others focus on disclosure, human review, data access, record retention, or discrimination. Their coverage thresholds and effective dates vary. Yet together they mark a transition from the first generation of workplace AI regulation, which concentrated largely on hiring transparency, to a more demanding era of governance over the entire employment lifecycle.

2. The First Generation: Bias Audits and Hiring Transparency

Agentic AI exposes the limits of the first generation approach, which is heavily reliant on bias audits as a method of preventing algorithmic discrimination. An annual audit may provide useful information about a defined screening tool at a particular point in time. But an AI agent may interact with candidates, alter the sequence of questions, synthesize multiple data sources, recommend next steps, and communicate outcomes. As the system takes on more functions, employers will need to ask whether an audit of one component adequately captures the risks created by the full workflow, including changes in configuration, data inputs, model versions, and human reliance.

Illinois illustrates the movement from tool-specific regulation to decision-focused regulation. The Artificial Intelligence Video Interview Act began with a comparatively narrow use case. It requires notice, an explanation of how the technology works and what characteristics it evaluates, and consent before an employer uses AI to analyze an applicant-submitted video interview. It also regulates sharing and deletion of the video.[49]

The 2026 amendments to the Illinois Human Rights Act take a broader approach. The statute now addresses the use of AI with respect to recruitment, hiring, promotion, discipline, discharge, training, tenure, and other terms or conditions of employment. It prohibits uses that have the effect of subjecting employees to unlawful discrimination and requires notice when AI is used for covered employment decisions.[50] The direction, however, is unmistakable. Illinois is no longer regulating only a particular interview technology. It is regulating AI-assisted employment decisions. The focus is shifting from what the product is called to what the system does and what outcome it helps produce.

3. The Second Generation: Governance, Explanation, and Human Review

Colorado's Automated Decision-Making Technology Act applies when automated decision-making technology materially influences a consequential decision, including an employment decision. For employers, the core obligations are procedural. A deployer must provide pre-use notice. After an adverse consequential decision, it must provide a plain-language description of the decision and the technology's role, along with instructions for obtaining more information. On request, the affected person may correct personal data and seek meaningful human review and reconsideration.[51]

Connecticut's new law adds another model. Its automated employment decision technology provisions require plain-language disclosure when a person is interacting with the technology, unless the interaction would be obvious to a reasonable person. Employers must also provide pre-use information about the technology's purpose, the nature of the decision, the categories and sources of personal data processed, and how the data will be assessed.[52] Connecticut also addresses accountability through its existing civil-rights framework. Under the state's antidiscrimination law, an employer's use of automated employment decision technology is not a defense to a discrimination complaint under state law. Evidence of anti-bias testing or similar proactive measures may be relevant, but the employer cannot simply point to the technology as the responsible actor.[53]

4. California Enters the "No Robo Boss" Era

California first made its accountability position explicit through regulations under the Fair Employment and Housing Act. The regulations apply to automated decision systems that make or facilitate decisions concerning an employment benefit. They confirm that using such a system or related selection criteria in a manner that discriminates against an applicant or employee may violate FEHA. They also treat anti-bias testing, or the absence of it, as potentially relevant evidence and require retention of covered employment records, including automated-decision-system data, for four years.[54]

The regulations matter because they reject the notion that automation creates a liability gap. If a computational process facilitates a discriminatory decision, the availability of a vendor, model, or algorithm as another causal actor *does not erase the employer's obligations under FEHA*. The employer still selected, configured, deployed, and relied on the process in an employment setting.

California moved beyond anti-discrimination principles with the No Robo Bosses Act. The law prohibits an employer from relying solely on an automated decision system to make a disciplinary or termination decision. When the employer primarily relies on the system, a human reviewer must corroborate the proposed decision using other information, such as performance evaluations, witness or customer interviews, or personnel records. If the recommendation cannot be corroborated, it cannot be used as the basis for the action.[55]

> "Despite their differences, these laws increasingly converge around one rule: employers cannot outsource accountability.

The distinction between "solely" and "primarily" is important. The law *does not ban* automated assistance in discipline or termination. It prohibits sole reliance and imposes independent corroboration when the system is the primary basis for the decision. That nuance should shape policies, training, and documentation. A manager who merely clicks "approve" has not supplied meaningful review. The human reviewer must *evaluate evidence outside the system's output* and must be prepared to reject the recommendation.[56]

5. The "Can't Blame AI" Principle

Despite their differences, these laws increasingly converge around one rule: employers cannot outsource accountability. In Illinois, an AI-assisted employment decision remains subject to the Illinois Human Rights Act.[57] In California, the FEHA applies when an automated system makes or facilitates a discriminatory employment decision, and the No Robo Bosses Act requires human corroboration in specified disciplinary and termination decisions.[58] Connecticut expressly provides that use of automated employment decision technology is not a defense to discrimination.[59] Colorado places notice, explanation, correction, and human-review duties on the deployer that uses the technology.[60]

That convergence has practical consequences for vendor management. A contract can allocate financial risk between an employer and a technology provider, but it cannot guarantee that a regulator, court, applicant, or employee will look only to the vendor. Nor should an employer assume that a vendor's generic representation that a product is "compliant"

answers the questions that matter: What data does the system use? What outcomes does it influence? Can the employer test for adverse impact? Can the employer preserve relevant records? Can a human understand and override the output? Will the vendor provide information needed for notices, access requests, investigations, and litigation?

The emerging "No Robo Boss" era is not a rejection of workplace technology. It is a rejection of unaccountable automation. Employers may delegate tasks to an AI agent, but they cannot delegate their obligations under anti-discrimination, privacy, transparency, and workforce-protection laws. The organizations best positioned to use agentic AI will not be those that remove people from every decision. They will be those that know where human judgment matters, preserve it deliberately, and can prove that it was exercised.

IV. Key Labor and Employment Law Considerations for Agentic AI

A. Legal Risks of Precursor AI

The precursor AI described in Section I – predictive AI, generative AI, NLP, and computer vision – have distinct objectives and outputs, but they expose employers to overlapping legal risks when used in connection with employment decisions such as recruiting, screening, hiring, managing, monitoring, compensating, promoting, disciplining, and terminating workers. Employment-related laws generally apply according to the employer's operations and personnel decisions, regardless of whether the employer used a person, software, or an automated system to carry out the operations and decisions. In addition, inaccuracies or other incorrect output due to user or system error may also give rise to legal risk.

1. AI-specific requirements

Employers should assess coverage of AI systems in use not only with respect to generally applicable employment laws but also to an expanding, jurisdiction-specific set of laws and regulations setting guardrails on how employers may use AI tools. An inventory of these laws within the United States will be discussed below. Multinational employers may also encounter the European Union AI Act's requirements for certain employment-related systems; these requirements will also be discussed in detail. Because coverage may turn on an AI system's function, whether the output underwent human review, the weight given to its output, the worker's or job's location, and the type of data processed, employers need coordinated legal, HR, procurement, privacy, and data security review before deployment and throughout the vendor relationship to ensure compliance with the laws currently in effect.

2. Discrimination and bias

State and federal antidiscrimination laws generally recognize two theories of liability for employment discrimination – disparate treatment and disparate impact – and AI implicates both.[61] Disparate treatment occurs when an employer fails to hire, discharges, or otherwise discriminates with respect to an individual's compensation, terms, conditions, or privileges of employment because of a protected characteristic.[62] Disparate treatment can arise in connection with the use of an AI tool when automated systems learn from biased training data to recognize and discriminate against protected characteristics without being explicitly programmed to do so.[63] Predictive AI tools can reproduce or amplify inequities embedded in historical hiring, promotion, performance, or compensation data. Even when a model excludes protected characteristics from screening, seemingly neutral variables such as ZIP code, schools attended, employment gaps, word choice, or wage history may operate as proxies, and unrepresentative training data may cause materially different error rates across groups.[64] NLP-based screening and interview tools may struggle with or misinterpret accents, speech differences, nonstandard syntax, culturally specific language, or disability-related communication patterns.[65] Computer vision systems may perform unevenly based on lighting, assistive devices, facial differences, or use of religious headwear, and generative AI may introduce stereotypes or exclusionary criteria into job advertisements, interview questions,

candidate summaries, and performance narratives.[66] In this way, algorithmic bias can affect who sees an opportunity, advances through resume screening, receives a workplace accommodation, or is selected for promotion, creating potential disparate treatment exposure.

Under the disparate impact theory of discrimination, plaintiffs may prove discrimination without proving intent to discriminate if an employment policy or practice disproportionately affects a protected group.[67] AI tools can magnify the risk of disparate impact claims as follows:

> In analyzing a large quantity of data, an algorithm might identify a statistical correlation between a specific characteristic of a job applicant and future job success that nevertheless lacks a causal relationship. Furthermore, employees and applicants also may have an easier path alleging class-wide discrimination claims if the employer uses the same AI tool or algorithm to assess an entire pool of candidates. Put differently, if an algorithm is applied across a group—or several groups—of applicants, the algorithm may easily provide the common questions of law or fact necessary to certify a class under Rule 23 of the Federal Rules of Civil Procedure.[68]

A vendor's assurance that a tool is "bias free," or a general audit performed across all users, may not establish that the system is valid or equitable for the employer's particular jobs, applicant population, configuration, and decision process. Human review is also not a complete safeguard when reviewers lack information about the model's data and limitations, defer reflexively to an algorithmic score, or cannot identify patterns visible only across many decisions.[69]

3. Privacy and data protection concerns

Use of AI systems may require or encourage collection of large volumes of applicant, employee, customer, and consumer information. Predictive and NLP tools may infer sensitive attributes from ordinary data; generative AI users may inadvertently place confidential personnel or medical information, attorney-client privileged or trade secret information, or proprietary customer information into prompts; and computer vision may capture facial geometry or other biometric identifiers.[70] Workplace monitoring can also collect personal data beyond what is necessary for the stated business purpose. Depending on the jurisdiction and data involved, employers may face notice, consent, data minimization, retention, security, access, deletion, and vendor-contracting obligations under a jurisdiction's general data privacy statute. Many U.S. states have data privacy laws in effect, and multinational employers must also comply with the EU's GDPR and other global privacy regulations. Biometric information warrants particular caution because some laws specifically regulate its collection, storage, disclosure, retention, and destruction, and a compromised biometric identifier cannot simply be replaced like a password.[71] For example, use of computer vision tools for facial recognition functions create risk for employers in Maryland and Portland, Oregon due to the laws in effect in those jurisdictions intended to protect workers' and consumers' privacy rights.[72]

4. Unfair competition and trade secret protection

Employees may expose an employer's trade secrets or confidential information belonging to a customer by entering source code, pricing, strategy, customer information, product plans, or other proprietary material into an AI system whose provider may retain prompts, use them to improve models, or permit access inconsistent with applicable confidentiality duties. This disclosure or use can support a misappropriation claim and may also undermine the employer's ability to show that it took the "reasonable measures" required to maintain trade secret protection.[73] Conversely, AI-generated competitive intelligence, code, or business content may reproduce or derive from information obtained under circumstances that restrict its use, creating risk when the employer cannot establish provenance. AI output used in recruiting, marketing, or commercial communications may also create unfair competition exposure if it falsely describes a product or service, causes confusion about source, sponsorship, or affiliation, or improperly uses a competitor's marks.[74]

5. Inaccuracy and hallucination

Predictive systems can produce false positives, false negatives, or unreliable forecasts when data are incomplete, unrepresentative, stale, or poorly matched to the workplace context.[75] NLP and computer vision tools can misclassify language, make speech transcription errors, or misidentify people or events, while generative AI may hallucinate facts, citations, qualifications, policy requirements, or explanations that sound plausible but are unsupported.[76] Relying on faulty outputs in connection with consequential decisionmaking may result in failure to follow legal regulations or policies, negligent or defamatory communications, and flawed investigations. The inability to explain and defend the decision opens an employer up to significant exposure. Use of false or misleading AI-generated information in court filings or in reports and other information submitted to regulatory bodies is also a major risk.

B. Agency Principles and the Allocation of Workplace Authority

The next phase of workplace AI is not defined by better tools, but by greater autonomy. When an employer intentionally confers autonomy on autonomous systems to perform workplace functions, existing agency principles may become increasingly important in determining how responsibility is allocated when those systems influence or execute employment-related decisions. The key question may be whether and to what extent the organization's decision to delegate autonomy to the AI alters traditional principles of accountability?

Under federal statutes such as Title VII, the ADA, and the ADEA, the term “employer” includes “any agent of” an employer.[77] However, the word “agent” now carries several meanings including an AI system in the technological sense described as “systems where LLMs dynamically direct their own processes and toll usage, maintaining control over how they accomplish tasks,”[78] an agent under common law doctrine, or an employer's statutory "agent" under federal anti-discrimination laws. Those concepts are not necessarily interchangeable. Employment law has nevertheless long tied employer responsibility to delegated authority, particularly where employers empower supervisors and other agents to act on their behalf. Agentic systems raise similar questions when organizations authorize autonomous systems to influence or execute workplace decisions.

1. Authority Conferred at Deployment

Two frameworks bear on whether an employer answers for an AI system's conduct: general agency principles, which turn on delegated authority and control, and California's ADS regulations, which expressly reach ADS-enabled employment decisions. Both point toward employer responsibility.

Delegated authority. The EEOC has long recognized that employers may act through agents and has defined an agent as an individual or entity having "the authority to act on behalf of, or at the direction of," the employer.[79] The agency relationship therefore turns on delegated authority rather than employment status alone. Consistent with that approach, liability questions involving agentic AI may focus less on whether an autonomous system qualifies as an employee or supervisor and more on the authority the employer conferred on the system.

Routes to principal liability. Restatement (Third) of Agency § 7.03 recognizes that a principal may be subject to liability where an agent acts with actual authority, where the principal ratifies conduct, where the principal negligently supervises or controls the agent, or where the principal delegates performance of a duty to another.[80] Although those principles developed for human actors, they focus on delegated authority and responsibility rather than solely on the characteristics of the actor performing the task.

Is an AI system an agent? Whether an AI system can itself be a common-law agent remains unresolved, but recent scholarship argues that it cannot. Some argue that the common-law definition of agency “does not plausibly encompass agentic AI” because the relationship requires two legal persons capable of participating in an ongoing, consensual fiduciary relationship. As such, an AI system “lacks the capacity to be an ‘agent’” and is instead an instrumentality created

and deployed by human actors.[81] Others have similarly observed that "[s]oftware cannot have a consensual relationship with a human. As such, agency law doesn't provide a rule to justify or explain the basis for the relationship between software and humans."[82]

Why control matters. Control matters as well, and that requirement is not satisfied merely because one party creates a mechanism through which another party performs a function. The U.S. Supreme Court's decision in *General Building Contractors Ass'n, Inc. v. Pennsylvania* illustrates why control is central. There, construction employers had agreed to an exclusive hiring hall administered by a union that engaged in intentional discrimination. Although the hiring hall existed through the employers' agreement and supplied workers to their common workforce, the Court rejected the conclusion that the union was therefore the employers' agent. "At the core of agency," the Court explained, is a fiduciary relationship arising from consent that another act on one's behalf and "subject to his control."[83] The relationship between the employers and the union could not be characterized as one of principal and agent because the employers lacked control over the union's discriminatory hiring practices. Agentic AI presents the inverse circumstance: the employer selects the system, defines its permissions, and determines the role of human oversight. A system that is not itself an agent may therefore be treated as the employer's instrument, with its conduct analyzed as the employer's own.

California's broader approach. California's recently enacted automated decision system (ADS) regulations under the Fair Employment and Housing Act take a broader approach to agency in the employment context. The regulations define an employer's "agent" to include "any person acting on behalf of an employer, directly or indirectly, to exercise a function traditionally exercised by the employer or any other FEHA-regulated activity," including when the activity or decision is conducted "in whole or in part through the use of an automated decision system."[84] An agent falling within that definition is also treated as an "employer" for purposes of FEHA, and employers and other covered entities are liable for discriminatory actions by their agents committed within the scope of the agency relationship.[85] These regulations do not appear to designate an ADS, standing alone, as a legal agent. Instead, they expressly reach persons acting for an employer in traditionally employer-controlled functions, including applicant screening, hiring, promotion, compensation, benefits, and leave, when those functions are performed in whole or in part through an ADS. Unlike the common law control test, California's regulatory framework expressly brings ADS-enabled employment activity within its definition of an employer's agent.

Ultimately, the system's conduct is likely to be attributed to the employer that deployed it. The practical question is what the system has been permitted to do and how it is presented to the workforce. If employees are directed to an agent for leave or accommodation information, for example, courts may consider whether workers would reasonably understand it to speak for the employer.

2. Respondeat Superior and the Supervisor Question

Agentic AI systems increasingly perform functions historically carried out by managers and supervisors. An agentic system may recommend discipline, assign work, evaluate performance, communicate with employees, or influence hiring decisions. Although the technology differs from a human supervisor, the underlying legal question remains familiar: when does conduct carried out on behalf of an employer become attributable to the employer itself? Two doctrines bear on that question: respondeat superior, which asks whether the conduct served the employer's objectives, and supervisor liability rules, which ask whether the actor was empowered to take tangible employment actions. Under both, the answer turns on the authority the employer conferred.

Scope of employment. Even if an AI system is not itself an employee or agent, the doctrine's rationale points the same way. Under the doctrine of respondeat superior, an employer may be held responsible for conduct undertaken by an employee acting within the scope of employment. As one court explained, "[a]n essential ingredient of agency is that the agent is doing something at the behest and for the benefit of the principal. It must be the affairs of the principal, and not solely the affairs of the agent, which are being furthered in order for the doctrine of respondeat superior to apply."[86] Because an AI system has no independent interests, its conduct will ordinarily be directed toward the objectives its deployer assigns.

Supervisor status under federal law. Under federal law, an employee is a supervisor "when the employer has empowered that employee to take tangible employment actions against the victim, i.e., to effect a 'significant change in employment status, such as hiring, firing, failing to promote, reassignment with significantly different responsibilities, or a decision causing significant change in benefits.'"[87] Supervisor status matters because the *Faragher/Ellerth* affirmative defense is unavailable where a supervisor's harassment culminates in a tangible employment action.[88] Although that framework arose in harassment cases, its test, whether the actor was empowered to take tangible employment actions, maps closely onto what agentic systems may be authorized to do. An employer that lets a system reject applicants or reassign work without human approval has conferred the very authority that defines supervisor status.

> "
>
> Emerging AI governance frameworks increasingly focus on the quality of human oversight rather than its existence alone.

New York City's broader rule. The attribution question carries particular significance under the New York City Human Rights Law (NYCHRL). The law establishes three routes to employer liability for discriminatory conduct by an employee or agent: strict liability where the wrongdoer exercised managerial or supervisory responsibility; liability based on actual knowledge and acquiescence; and liability based on constructive knowledge coupled with a failure to exercise reasonable diligence.[89]

The open question. Courts have not yet resolved whether the functional autonomy conferred on an AI system is sufficient to implicate the supervisory liability framework under the NYCHRL.[90] On one hand, an AI system may be authorized to perform functions traditionally associated with supervisors, such as directing employee conduct, communicating disciplinary decisions, or acting with the apparent authority of human resources. On the other, employers may argue that AI systems remain tools operating within a structure of human oversight and therefore do not independently possess supervisory authority.

3. Human Oversight and the Decision-Maker

Employers have long relied on human oversight as a safeguard, and employment law has long looked for a decision-maker whose reasons can be examined. Keeping a "human in the loop" is standard guidance; however, agentic AI complicates both because human involvement exists on a spectrum. A person may independently evaluate the relevant information, approve an AI-generated recommendation, monitor many automated decisions at once, or intervene only after the agent has acted.

Emerging AI governance frameworks increasingly focus on the quality of human oversight rather than its existence alone. NIST's AI Risk Management Framework emphasizes that human roles and responsibilities in overseeing AI systems should be "clearly defined and differentiated" and recognizes that certain forms of human-AI interaction may amplify, rather than mitigate, decision-making risks.[91] The resulting inquiry may therefore focus not simply on whether a human remained somewhere in the process, but on whether that individual exercised independent judgment and retained meaningful responsibility for the ultimate decision.

For employers, the practical lesson is that the value of human oversight depends on where it sits and what the reviewer can do. Oversight placed before consequential actions, by someone with the information and authority to change the outcome, is more likely to function as the safeguard employers intend.

Agency law has long focused on authority, delegation, supervision, and accountability. Agentic AI increases their significance. Courts and regulators may focus on the authority granted to the system at deployment, the foreseeability of the resulting conduct, and the individuals or entities that retained the power to prevent the outcome. Employers may delegate workplace functions, but the allocation of responsibility for their consequences will likely follow the authority they conferred.

C. Dynamism and Discrimination Risk

Agentic AI's dynamism, the way each decision shapes the next, can turn a facially neutral objective into a discriminatory pattern that neither the employer nor the developer intended or can readily explain. That combination may make disparate impact claims easier to prove and harder to defend. Consider an autonomous vehicle: faced with a fork in the road, the algorithm must decide whether to take the right or left fork. But if the vehicle turns left, the algorithm sees a different traffic pattern than if the vehicle had turned right, and all its subsequent decisions reflect the consequences of that prior choice.

The ability of AI to discern an optimal decision based on ever-changing information creates unique and difficult challenges for employers and developers. Consider the following hypothetical:

- An agent is tasked with developing a workforce that is both highly efficient and likely to include loyal and long-serving employees.
- The agent observes that when employees are teamed with others who reside in similar neighborhoods, and attend the same or similar schools, their turnover rates are lower and production is greater. We can call this trait "connectedness," which may be a neutral trait, but its impact may be less benign. For example, these traits may be correlated with ethnicity, and result in members of a given ethnic group preferred for some jobs but not others. With each hire, the connectedness within the group increases because of the agent's search for similarity, and any disharmony associated with a "non-traditional" hire may impact an ever-greater number of employees.
- As a result, the agent steers a path that tends to segregate the workforce. Note, however, that the agent has no information regarding any demographic characteristic but has learned on its own that group "connectedness" is an important aspect of productivity, and seeks out hires who are most likely to connect with those it previously slotted into particular positions.
- Managers may recognize this outcome, but neither they nor the developer of the tool may have an inkling of how this came about, let alone how to remedy it. The agent was tasked with maximizing productivity and found a relationship in the data which it exploited in pursuit of efficiency. But its sequential decisions expose the employer, and perhaps the developer, to potential liability.

Burden of proof. Title VII of the Civil Right Act of 1964, as amended, prohibits selection methods that are neutral on their face but result in decisions that disproportionately impact members of a particular demographic group.[92] It is the plaintiff's burden to identify the particular facet in a multi-faceted process that causes the disparity, unless that hurdle cannot be disentangled from other aspects of the decision. In that case, the entire process may be analyzed as an integrated whole. Thus, if the value of "connectedness" overrides the influence of other selection criteria developed by the agent, the disparity might be evident at the "bottom line," *i.e.*, the grouping of employees into ethnic enclaves, which would satisfy the plaintiff's burden of proof.

Establishing a defense. The employer would seek to avoid liability by demonstrating the ethnic groupings are a by-product of selection criteria that are job-related and consistent with business necessity, but it may have no means of identifying the criteria the algorithm developed that yielded this pattern of segregation.[93] Were the actual selection criteria known to the employer, along with the tendency of employees who shared the same background to form more productive groups, the employer in fact may have a defensible case. However, the rationale for an algorithm's decisions may be no more transparent than the reason an autonomous vehicle turned right on First Street rather than Second Street, when traffic on both streets seemed similar. If the opacity of the algorithm's decisionmaking process prevents an employer from being able to demonstrate that its selection policy is job-related for the position in question and consistent with business necessity, the employer cannot defeat a disparate impact claim.[94]

D. The Ability to Access Data Across Systems and Data Privacy

The characteristic that makes agentic AI particularly useful, its ability to access, connect, and act on information across systems, creates data privacy challenges. Rather than treating an AI agent as simply another software application, employers may need to consider the data the agent can reach, the information it creates, what it can do with that information, and who ultimately receives it.

1. Configuration Conundrums: When AI Can Find Everything

Many privacy issues surrounding agentic AI begin with a deceptively simple question: What information can the agent access? Organizations have long relied on permissions, data classification, network segmentation, and other controls to determine who can access sensitive information. Agentic AI can increase the importance of those controls because an agent may retrieve information from multiple repositories far more efficiently than an individual employee searching those repositories manually. Consider an incorrectly permissioned folder containing employee medical information. An employee might historically have had technical access to the folder without knowing that it existed. An AI agent capable of searching broadly across organizational data may make the information significantly easier to surface.

Fundamental questions. Employers considering agentic AI therefore may want to revisit several fundamentals. Is the information available to an agent current and accurate? Are obsolete policies and duplicate personnel records appropriately archived? Is information classified correctly? Most importantly, does the agent inherit appropriate user permissions, so that an employee cannot use the technology to obtain information the employee should not otherwise access?

Security boundaries. These issues extend beyond confidentiality to cybersecurity. Connecting systems that previously operated separately can affect existing security boundaries. A compromised agent or credential potentially provides access to multiple sources rather than one application. Organizations might therefore consider controls such as least-privilege access, enhanced logging, anomaly detection, and periodic access reviews.

Access and visibility. This is one area where AI itself may offer part of the solution. Security tools increasingly can analyze activity patterns to identify anomalous behavior. The broader point is that deploying agentic AI may call for considering both sides of the equation: increasing an agent's access may warrant increasing visibility into what the agent is doing with that access.

2. Agentic AI Can Create a New Category of Employee Data

Privacy considerations do not end with information an employer already possesses. Agentic AI may generate additional information about employees. For example, an agent might learn an employee's writing patterns, preferred applications, routine working practices, voice, appearance, or communication preferences. With appropriate capabilities, that information could allow an agent not merely to assist an employee but to emulate aspects of the employee, such as drafting correspondence in the employee's style or generating audio or video communications. That possibility introduces several overlapping legal questions.

Biometric information. Voice or face-based functionality should receive particular attention. Illinois' Biometric Information Privacy Act (BIPA), for example, expressly covers voiceprints and scans of face geometry and imposes requirements concerning notice, written release, retention, disclosure, and safeguarding when its provisions apply.[95]

Identity and publicity rights. AI-generated likenesses also may implicate state identity, likeness, or right-of-publicity laws. If an organization proposes to have an employee's digital likeness, cloned voice, or avatar communicate externally, employers may want to evaluate whether authorization is required and clearly establish the purposes and limits of that authorization. There is also a practical security dimension: technology sufficiently sophisticated to reproduce an

employee's voice or appearance potentially creates a valuable impersonation credential. Employers might consider whether such profiles warrant heightened access controls, restrictions on reuse, and specific deletion requirements when employment ends.

Individual privacy rights. Comprehensive data protection laws add another layer. In the United States, only the California Consumer Privacy Act applies to the personal data of the workforce, but virtually every other highly developed country has data protection laws that protect such data.[96] These laws provide rights including access, correction, and deletion, subject to statutory exceptions. Employers deploying AI-generated employee profiles should consequently consider whether existing processes for responding to privacy rights requests can locate and act upon information residing inside the AI ecosystem, rather than only information in traditional HR databases.

3. Notice and Consent: Who Knows the AI Is There?

Agentic AI also raises questions about notice, consent, and transparency. First, existing employee privacy notices may not adequately describe the new information being collected or new purposes for which existing information will be used. An organization that begins creating employee voice profiles, for example, may need to analyze whether applicable notices sufficiently describe that processing.

A second question is whether people interacting with an AI agent know that it is AI. This issue can arise internally, when one employee interacts with another employee's agent, and externally, when an agent communicates with customers, applicants, vendors, or other third parties. As a growing number of laws require notice to individuals who might not otherwise realize that they are interacting with AI,[97] organizations may want to establish a baseline rule for when agents identify themselves rather than treating disclosure as a system-by-system decision.

Recording and communications laws deserve particular attention. An agent that listens to a telephone call or virtual meeting in order to transcribe, summarize, or take action may implicate federal and state wiretap laws. Requirements differ significantly by jurisdiction and circumstances. Recent litigation involving AI meeting assistants in which the plaintiffs allege that an AI notetaking service recorded and transcribed conversations without appropriate consent shows how the issue can arise.[98] The case illustrates the potential importance of whether a technology provider acts solely on behalf of the customer or uses intercepted information for its own purposes. If the service provider that supports the AI agent uses the data collected by the agent for its own purposes, even in part, the service provider could be characterized as a third party "listening in" on employees' conversations. Without clear consent, the employer potentially could be liable for aiding and abetting the service provider to violate wiretap laws. For employers, this highlights the importance of vendor contracting. Agreements can specify whether the provider acts only on the employer's instructions and whether conversation data may be retained, reused, or used to train the vendor's models.

4. When the Agent Crosses Borders

An AI agent may also move or expose information across national boundaries simply because its underlying infrastructure, support personnel, or sub-processors are located elsewhere. For multinational employers, the GDPR and other international data-transfer requirements impose familiar requirements. Although the U.S. has historically placed few restrictions on the transfer of personal data overseas, the U.S. appears to be moving toward more protections. In particular, the Department of Justice's Data Security Program restricts or prohibits certain transactions that provide countries of concern or covered persons access to specified categories of bulk sensitive personal data, including Social Security numbers, biometric, precise geolocation, health, and financial data.[99]

This does not turn every overseas AI deployment into a prohibited transfer. It does, however, make knowing where an agent's data goes and who can access it increasingly important. The DOJ itself advises organizations to "know their data," including its types, volumes, uses, and relevant transactions. Employers evaluating an agentic AI vendor may accordingly want visibility into hosting locations, sub-processors, remote support access, and contractual restrictions on onward access.[100]

5. Vendor Management Becomes AI Governance

Finally, agentic AI can make conventional vendor diligence considerably more consequential. A provider may hold not only employer data but also the credentials or integrations necessary to act within employer systems. Contract review therefore can extend beyond standard confidentiality language. Depending on how an AI system will be deployed, a vendor contract might address security requirements, access controls, permitted use purposes, model training restrictions, data retention and deletion, sub-processors, security breach investigation and notification, audit and logging capabilities, allocation of incident cooperation, data location, allocation of responsibility for security vulnerabilities, service levels, indemnification, and procedures for terminating the agent's access.

The central privacy challenge posed by agentic AI is therefore not entirely new. Organizations already manage access permissions, privacy notices, vendor relationships, biometric information, security, recording laws, and international transfers. Agentic AI brings these familiar issues together, at scale and sometimes in a single system.

E. Optimizing Productivity, Labor Budget, and Wage and Hour Compliance

Agentic AI has the potential to transform one of the most troublesome tasks facing employers with large hourly workforces: deciding which workers to schedule at which shifts to meet production needs and keep overtime costs low while maintaining compliance with wage and hour laws. An agentic scheduling system might be asked to build next month's schedule based on expected customer demand, employee availability, skill sets and approved leaves, and labor budgets. More sophisticated systems may continue working after the schedule is published, responding to call-offs, filling open shifts, and modifying staffing as business conditions change.

1. Delegate the Work and Mitigate the Risk

The appeal of using AI in this area is easy to understand. Scheduling can be complicated, particularly in retail, hospitality, and healthcare environments in which demand can fluctuate vastly day to day, even hour by hour. Managers must balance numerous competing factors simultaneously. But the same autonomy that makes agentic scheduling powerful presents important wage and hour considerations. A scheduling agent is intended to do more than just make recommendations to a manager: it may be permitted to actually change an employee's schedule, even send them home early, communicate with the employee about schedule changes, and reassign work tasks. Complying with wage and hour obligations may often be at odds with the agent's main objective: to keep the labor budget down while ensuring sufficient people power to achieve production or sales goals. The same constraints placed on managers do not disappear when the decisions are delegated to AI.[101]

2. Scheduling Compliance Is More Than Just Overtime Rules

An employer deploying an autonomous scheduling tool may naturally want to instruct the system to minimize overtime while maintaining appropriate staffing levels. But incorporating overtime rules is only one consideration. Depending on the jurisdictions and workforce involved, a scheduling agent may need to also address:

- Multistate meal and rest period compliance, including whether the staffing levels will permit the employees to realistically take those breaks in a compliant manner;
- Right to rest and "clopening" protections, split shift and spread-of-hours rules, and "day of rest" requirements; and
- Collective bargaining agreements that may impose additional requirements concerning scheduling, shift premiums, overtime distribution, seniority, or schedule changes.

A schedule that looks compliant on paper may collapse if it is impossible to follow or when unexpected customer demand causes the agent to reallocate available workers to a different location or department.

Last-minute and real-time adjustments. Agentic systems add another dimension because a schedule may no longer be static. The very feature that makes an autonomous agent useful may be its ability to monitor business conditions and continually optimize staffing. A tool might respond to unexpectedly low demand by delaying a shift, telling an employee not to report, or sending an employee home early, all of which could trigger reporting time pay or guaranteed minimums under a collective bargaining agreement. An agent might respond to high demand by extending shifts or asking employees to report on short notice even if they have already worked that day, implicating potential split shift or reporting time penalties. Further last-minute changes can implicate predictive scheduling laws and predictability pay or alternative workweek schedule penalties. For employers operating in multiple jurisdictions, the challenge may be particularly significant because a scheduling strategy permissible at one location may trigger additional obligations elsewhere.

An agent may unwittingly engage an employee in work. Autonomous communication presents an additional wage and hour risk. An agent charged with maintaining staffing levels might consider it reasonable to contact employees outside their scheduled hours to confirm availability, respond to shift change requests, solicit volunteers for open shifts, or even ask employees to complete a preliminary task. Each individual interaction might appear insignificant but arguably may trigger compensable time in certain jurisdictions. Agentic technology can make that issue more consequential because the system may itself create compensable work.

Agentic AI can have recordkeeping implications. The increased volume of data generated by agentic systems may also make recordkeeping more complicated in some respects. Time punches may no longer be the only source reflecting when work occurs. Scheduling platforms, task assignment agents, location information, agent communications, and system logs may each generate data that is relevant to when an employee is performing work. When these sources of information are inconsistent with an employee's time punches, this can lead to claims that an employee has worked off the clock.

Agentic AI may change the jobs of exempt employees. The wage and hour implications of agentic AI are not limited to nonexempt employees. Many of the core functions that agentic AI may be asked to perform have until now been performed by an exempt manager. Managers who historically decided how many employees to schedule, which employees possessed the appropriate skills, how to respond to fluctuations in demand, whether to call in additional employees for a shift, how to accommodate time off requests, and how to control overtime may be freed up to perform department goal setting or focus on process improvement. This redeployment of manager focus will require reevaluating managers' exemption status to determine whether scheduling duties are replaced with equally exempt duties, and that the manager continues to exercise the same level of independent judgment and discretion necessary to satisfy the exemption on which the employer relies.

Wage and hour audits may be even more important. Despite these risks, agentic scheduling can be valuable and beneficial. This technology is ultimately likely to provide greater opportunities to incorporate compliance into scheduling and workforce management processes with the right guardrails in place prior to deployment, including identification of all wage and hour rules that apply to the relevant workforce and testing of the system against these constraints. Particular attention should be paid to jurisdiction-specific requirements such as state and local rules and collective bargaining agreements that a generically configured scheduling tool is unlikely to be built around.

Finally, auditing should not end when the technology goes live. Agentic systems are expected to respond dynamically to changing conditions. Therefore, as changes are made to how schedules are created and how work is assigned, a system's deployment design must also adapt to ensure that the efficiencies AI produces do not result in unintended costs.

F. How Agentic AI Shifts Labor-Management Relations

Agentic AI is a new tool in the modern workforce, but far from the first technological advancement that has changed the nature of work. Whether it was the introduction of early stage robotics, or the onset of social media at work, the Board has applied traditional labor law principles and assessed the new workplace facts and realities within those overarching principles. What is somewhat different today with agentic AI is that many existing doctrines assume identifiable human

actors who assign work, exercise judgment, supervise employees, and communicate management decisions. Agentic AI may perform or influence portions of those functions with less direct human involvement and employers need to consider how that impacts the application of existing labor law principles.

1. Section 7 Rights in Every Workplace (Sections 7 and 8(a)(1))

Section 7 of the NLRA protects covered employees who engage in concerted activity for collective bargaining or other mutual aid or protection, regardless of whether a union represents them.[102] Protected activity may include employees acting together, bringing group concerns to management, or seeking to initiate or prepare for group action concerning workplace issues. Surveillance of protected activity, or creating the impression of surveillance, may violate Section 8(a)(1), while observing open protected activity generally is lawful unless the observation is "out of the ordinary."[103]

New subjects for concerted activity. Employees today may join together to oppose AI-driven monitoring, productivity expectations, scheduling practices, performance evaluations, and to raise job security concerns driven by AI. Employee discussions about an agent's effects on jobs, schedules, or productivity tracking may themselves be protected when they relate to terms and conditions of employment. While the topics may be new, the governing legal framework remains the same.

2. Organizing Risk and Employee Voice

Communication and transparency. Employees may experience ongoing adjustments to and questions about the future of their work driven by agentic AI advancements. They want to know what's coming and how it will affect them. Employees want to know if AI is monitoring and evaluating their work, whether it will replace them, and whether anyone cares what they think. At the same time, employers may not be able to predict how jobs, reporting structures, skill requirements, or career paths will evolve as systems gain capabilities or employers expand use cases.

Employers may hesitate to communicate before they have reliable answers, but delayed communication can create information gaps that allow others, including unions, to define the narrative surrounding AI adoption. Transparency need not require predictions about outcomes. It can cover process issues such as when and how information will be shared and employee opportunities for feedback. Transparency can also include answers to questions about what a tool does and does not do when an employer is ready to communicate that information.

White collar organizing. Agentic AI reaches directly into professional work. Engineers, accountants, analysts, and other knowledge workers may confront questions regarding job security, career development, skill relevance, and workplace influence. Employee engagement and communications plans should include messaging about when and how AI will affect professional judgment, development opportunities, and future career paths. Employers evaluating AI-related organizing risks should also include consideration of potential activity targeting professional and technical employees in addition to hourly operations populations.

3. Which Employees Are Included in a Bargaining Unit

Supervisory status. If organizing activity is underway, it is important for an employer to consider who may be a part of a potential bargaining unit. Under the NLRA, supervisory, managerial and confidential employees are not part of a bargaining unit. The law also places restrictions on what managers and supervisors can and cannot say and do during organizing. For example, the law prohibits an employer from threatening employees with adverse consequences if they support unionization efforts, and an employer may be held liable for statements made by its managers and supervisors. Section 2(11) of the Act defines a supervisor as an individual who holds specified authority in the employer's interest, including authority to assign, responsibly direct, discipline, or effectively recommend such action, when exercising that authority is not merely routine or clerical but requires independent judgment. [104]AI systems may take over some of these duties, scheduling employees, assigning work, monitoring productivity, identifying employees for coaching, or recommending discipline. When employers eliminate supervisory layers or shift authority or discretion from individuals to systems, the actual exercise of authority may no longer align with job descriptions, organizational charts, or historical classification decisions.

Bargaining unit composition. Under the Board's unit determination framework, employees in a petitioned-for unit must be readily identifiable as a group and share a community of interest. That analysis considers the employees' interests and working conditions rather than relying solely on job titles or organizational labels. Agentic AI may change the facts underlying that analysis by redistributing work, increasing functional integration, creating or eliminating classifications, changing supervisory structures, and/or blurring distinctions among analytical, technical, professional, clerical, and operational roles. Employers redesigning work around agentic systems may therefore need to consider whether historical classification and unit structures continue to match the organization's actual operations.

4. Duty to Bargain and Unilateral Change (Sections 8(a)(5) and 8(d))

An employer generally may not change a mandatory subject of bargaining (wages, hours, and other terms and conditions of employment, such as workloads, metrics, and staffing) without giving the union notice and a meaningful opportunity to bargain, unless the union has waived that right or the contract authorizes the change.[105] Some decisions require bargaining over the decision itself; others, only over effects (for example, whether employees will be provided with severance or retraining). Replacing unit work (work performed by bargaining unit employees) largely for labor cost reasons is a mandatory subject, while decisions about automation appear to fit into another category of decisions that involve the "scope and direction" of the business (known as *First National Maintenance* category III decisions)[106] are subject to a balancing test, with effects bargaining generally still required. For technology, the Board has asked whether a change is one of degree or of kind. The Board has also treated certain monitoring technology, such as hidden cameras, as a mandatory subject.[107]

Contract language that did not anticipate agents. Employers should also consider whether they have existing contract language that gives them the right to implement agentic AI unilaterally, and/or whether it may be prudent to bargain for language giving them future flexibility. Management-rights, technology, discipline, evaluation, and work-assignment provisions are read together, and many assume a human decision-maker, not a tool that acts, adapts, and is updated by a vendor. Moving unit work to an agent may also implicate work preservation or subcontracting clauses.

Notice of configuration and model changes. Because agents can be reconfigured, updated, or given more autonomy after deployment, unions may seek notice not only at adoption but also when an agent's configuration, underlying model, or scope of authority changes. Some parties may prefer a process, such as notice of material configuration changes and an opportunity to discuss them, over language tied to today's tools.

Common proposals. AI has, predictably, been a hot issue at the bargaining table. While unions have taken varied positions, union proposals have clustered in their focus on the same concepts. Specifically, unions have sought notice, information regarding the particular tools envisioned, human review of outputs, and commitments to training, work preservation, reskilling, redeployment, transition assistance, and joint committees. Because AI tools evolve and develop rapidly, employers should not focus only on presently existing tools, but in negotiating terms of implementation for all future agentic AI that the employer may want to implement during the contract term.

5. Information Requests and Employee Representation Rights (Section 8(a)(5) and Weingarten)

Information requests. An employer must provide information relevant and necessary to the union's role as bargaining representative. Information about unit employees' terms and conditions of employment is presumptively relevant; otherwise, the union must show relevance. An employer may raise legitimate confidentiality interests but generally must do so timely and seek an accommodation.

Unions may use this mechanism to seek information about an agent's configuration, decision criteria, logs, error rates, bias testing, and vendor documentation. Some of this data may be hard to produce: data may sit with a vendor, logs may not be retained in usable form, and a particular output may be difficult to explain fully. Employers may also have legitimate proprietary concerns, particularly with vendor information, which brings the confidentiality-and-accommodation framework into play. Employers may want to review vendor agreements for data access, retention, and confidentiality proactively in advance of any union request.

***Weingarten* rights.** A union represented employee may request a union representative at an investigatory interview the employee reasonably believes could lead to discipline. If an agent is used to gather an employee's account of an incident, for example through a chat interface, *Weingarten* rights could potentially attach depending on: whether the exchange is an "interview," whether those rights attach when the interview is conducted by a non-human, how a representation request made to an agent is recognized and honored, and what role the representative plays.

6. Discipline, Grievances, and Arbitration (Just Cause and Arbitrability)

When determining whether an employer had just cause for its action, arbitrators review whether the investigation was fair and objective and whether the employer presented sufficient evidence for its decision.

Explaining AI-informed discipline. When an agent informs discipline, the decision-maker must be prepared to explain the facts and reasoning for the decision. An answer that "the agent said that should be the result," or something similar is unlikely to satisfy a just cause inquiry. Employers should prepare to discuss how the agent is configured, how the tool was used, what the decision-maker reviewed, and how they used their judgment to make the final disciplinary decision.

Agent-generated evidence. An agent may collect evidence, evaluate it, and recommend or initiate action, all of which generates evidence that may be relevant in the grievance process and/or the arbitration hearing. Employers should collect this evidence and be prepared to reconstruct the agent's remit and actions. Of note, agents are often modified or updated, so this evidence must be collected contemporaneously.

The Federal Rules of Evidence do not bind arbitrators but offer a reference point: Rule 901(b)(9) permits authenticating a system's output with evidence describing the system and showing that it produces an accurate result, and proposed Rule 707, still under committee study, would apply expert-reliability standards to machine-generated evidence offered without an expert. While arbitrators and courts sort out how to handle AI generated evidence, employers should err on the side of over-documenting the tools they are using, the instructions given to the tools, the outputs, and ensure that all outputs are reviewed so that the final step involves the application of human judgment.

G. How Agentic Autonomy Shapes Workforce Restructuring and Reductions in Force

A traditional reduction in force is a discrete event. Management identifies a business reason, selects the affected positions, chooses the people in them, and issues notices on a known timeline. Agentic AI loosens each of those anchors. An agent can be the reason for a restructuring, by absorbing work and hollowing out roles, and it can also be the tool used to carry one out, by identifying redundancies, scoring employees, drawing on monitoring data, and drafting separation paperwork. Applicable laws ask who made the decision, when, and why. Agentic AI can make those questions harder to answer.

Not every agentic change ends in a separation. Many employers will redesign roles so that employees supervise agents rather than perform the underlying work, which raises questions about whether the new duties still satisfy the requirements for a FLSA overtime exemption (see Section IV.E), how performance standards are reset, and whether the new role's standards are job-related and accommodate disabilities (see Section IV.H).

1. Agentic AI as a Reason for Restructuring

Employers should expect increased scrutiny of the AI rationale for reductions, based both on worker protections and an interest in forecasting labor market trends.[108]

Federal WARN Act. The federal Worker Adjustment and Retraining Notification (WARN) Act requires covered employers (generally those with 100 or more employees) to give 60 days' notice before a plant closing or mass layoff.[109] Employment losses are aggregated over a 90-day period, so a series of individually small reductions may together trigger the statute.[110] That matters for agentic deployments, where displacement may occur through incremental role elimination

rather than a single decision, if the employer has not engaged in a pre-planned restructuring scheme that occurs in stages. Without tracking cumulative employment losses, an employer may cross a threshold without having decided to conduct a layoff.

State Mini-WARN Acts. Newer laws add a disclosure layer. As discussed in Section II.B, California will require Cal-WARN notices to include additional information when a layoff is caused in whole or in substantial part by AI or other automated technology, beginning January 1, 2027,[111] and Connecticut and New York require WARN-covered employers to disclose whether AI or other technological change is related to the layoffs. These laws raise a causation question the statutes do not fully answer: when AI is one of several drivers of a restructuring, how much is "substantial"? Whatever answer the employer reaches will appear in a notice that regulators, employees, and plaintiffs can read. It is advisable to be consistent with what the employer says elsewhere, including public statements about an "AI-first" strategy, investor communications, and internal messaging. Shifting or inconsistent explanations for an adverse action can themselves be evidence of pretext or worse.[112]

Bargaining Obligations and Concerted Activity. In unionized workplaces, a decision to replace bargaining-unit work to reduce labor costs is presumptively a mandatory subject of bargaining,[113] whereas a decision at the core of entrepreneurial control requires bargaining only over its effects on employees, not over the decision itself. (see Section IV.F.4).[114]

2. Agentic AI to Support Selection Decisions

Where an agent helps decide who is affected, familiar discrimination theories apply, and Section IV.C explains why agentic dynamism can make them harder to defend. Disparate impact is available under Title VII[115] and under the ADEA, where the employer bears the burden on the "reasonable factors other than age" defense,[116] and the EEOC's regulation asks whether the employer assessed adverse impact and took steps to reduce it.[117] Federal enforcement of disparate impact has narrowed, but private plaintiffs and state agencies still pursue that theory. California's automated decision regulations and Colorado's automated decision-making law reach AI-assisted termination decisions,[118] and under the California rules evidence of anti-bias testing, or its absence, is relevant to claims and defenses (see Section III.B).[119] Further, while other states' antidiscrimination statutes do not expressly address the use of AI in employment decision-making like the laws in California, Connecticut, and New Jersey, this does not eliminate the possibility of exposure to liability under other states' laws. Plaintiffs may pursue an algorithmic discrimination claim under a traditional disparate treatment or disparate impact theory of liability.

Absence and performance. The most obvious risk is that agent inputs can treat absence as poor performance. Inputs such as output volume, throughput, tool usage or "AI adoption" metrics, and sponsorship signals (i.e., indicators that supportive of AI) accumulate with presence, so an employee on protected leave, or with reduced capacity because of a disability or pregnancy, scores lower unless the system neutralizes those leave and accommodation periods. Failing to do so can violate the FMLA's bar on using leave as a negative factor,[120] and parallel state laws,[121] can support ADA, Pregnancy Discrimination Act, and Pregnant Workers Fairness Act claims.[122] An employee on leave may lawfully be selected for layoff for a legitimate reason, but the employer must show the employee would have been selected regardless of the leave.[123] A manager's warning that taking leave will lead to selection is itself interference.[124] Where inputs derive from workplace monitoring, the data may also capture medical information, implicating data privacy and confidentiality concerns (see Section IV.D).[125]

Selection criteria. Under the Older Workers Benefit Protection Act (OWBPA), an employer seeking waivers of ADEA claims from employees aged 40 or older in connection with a group termination program must provide understandable written disclosures identifying the covered group, eligibility factors, applicable time limits, and required job-title and age information.[126] When an agent helps determine participation or selection, the disclosures must accurately reflect the actual process, including relevant agent-applied factors, rather than describe a human-only process that did not occur. OWBPA regulations prohibit misleading waiver language and recognize that material mistakes, omissions, or misstatements may undermine a knowing and voluntary waiver.[127] Failure to satisfy OWBPA's requirements can render the ADEA waiver unenforceable.[128]

> The defensible restructuring is the documented one.

Human review and vendor liability. The lack of human review of AI-enabled functions, or merely nominal human review, may not protect the employer: one federal court has allowed claims against an AI vendor on the theory that it acted as the employer's agent (see Section IV.B),[129] and California and Connecticut bar a defendant from arguing that AI autonomously caused the harm.[130]

Identifying issues. Privileged adverse impact analysis, together with individualized review of employees on leave or with accommodations, before selection is final lets the employer correct problems early.[131]

3. Records, Preservation, Vendors, and Cross-Border Exposure

The defensible restructuring is the documented one. A contemporaneous record of the business reasons for automating a function, the selection criteria, the testing performed, and the human review actually done is the employer's principal answer to claims of pretext, insufficient disclosure, or other employment-related claims. The same record will be sought in discovery, along with agent instructions, model and configuration changes, scoring outputs, dashboards, and monitoring data, potentially from both employers and vendors (see Section V.C.). Plaintiffs may also ask a court to order preservation of those materials, and to order an independent audit of the selection inputs, weights, and outputs, before separations take effect.

H. Leave and Accommodation with Agentic AI

Leave of absence and accommodation law protects employees who need time away from work, or require changes to how work is done, because of a medical condition, pregnancy, family need, military service obligation or other protected basis. The framework is familiar. The Family and Medical Leave Act (FMLA), the Americans with Disabilities Act (ADA), the Pregnant Workers Fairness Act (PWFA), the Uniformed Services Employment and Re-employment Rights Act (USERRA) and a growing number of state and local counterparts require employers to respond to requests, determine eligibility, engage in an interactive process, keep medical information confidential, reinstate employees to active work after job-protected leave, and refrain from using protected leave or accommodation as a negative factor in employment decisions.[132] Those duties turn on what the employer knows, when the employer knows it, and how and when it responds, not on whether the person or system that received a request recognized it as one.

Agentic AI may dramatically change how those duties are carried out in three ways. Employers or third party administrators may deploy agents as the first point of contact to receive and process requests (see Section I.C.4), and agents may be treated as speaking for the employer (see Section IV.B.1). Agents and AI tools can also provide accommodation themselves by removing workplace barriers. And AI tools that employees are required to use, or that measure or monitor their work, may create new barriers to equal employment. Section IV.G.2 addresses how automated performance inputs may improperly penalize protected absence when employers select employees for a reduction in force. This section addresses day-to-day administration.

1. Agentic AI as Support for Leave and Accommodation Administration

Intake and notice. An employee need not mention the FMLA, USERRA or another statute by name, or use the word "accommodation"; communicating generally about a protected need can be enough.[133] An agent that handles HR requests may be the first to "hear" it, and the employer is likely to be charged with knowing what the agent receives. Deadlines also run from the employer's notice: an eligibility notice under the FMLA, for example, is generally due within

five business days.[134] An agent that gives an inaccurate eligibility answer, or steers an employee away from taking leave, risks interfering with the employee's rights.[135] Whether a protected disclosure reaches a trained human, and what the record shows as to when the employer received notice, may become contested questions.

Administration. Automated attendance, scheduling, and time-tracking tools may improperly count protected absences against employees or override approved accommodations, such as a modified schedule or remote work.[136] Intermittent leave tracking, leave year calculations, and restoration to the same or an equivalent position after leave are also areas where automation can produce errors.[137]

Confidentiality. Medical information must be kept confidential and separate from ordinary personnel files under the ADA and various other laws.[138] Agents with access to information across systems, or that retain memory of prior conversations, may carry it into tools that were not meant to hold it or provide access to persons that should not have visibility into it (see Section IV.D).

Evaluation. Productivity, throughput, and AI-usage metrics may need adjustment depending on the nature of an employee's approved leave or accommodation. The analysis in Section IV.G.2 of how such inputs treat absence applies equally to ongoing ratings, bonuses, and promotions. Accommodations may also legitimately change how an individual's output is measured, which raises the question whether an automated rating properly accounts for employees' protected absences or accommodations.[139]

2. Agentic AI as a Reasonable Accommodation

AI tools can themselves provide assistive accommodations, including transcription and note-taking tools, scheduling and task-management agents, and tools that support employees with cognitive, visual, hearing, or executive-function needs. An employer must provide a qualified individual with a reasonable accommodation that is effective, although not necessarily the one the employee prefers, unless it would cause undue hardship.[140] As such tools become cheaper and more common, undue hardship arguments may become harder to sustain.

Essential functions. When agents take over tasks, the essential functions of a job may change.[141] An agent's ability to perform a task does not by itself remove the need to accommodate an employee or reinstate the employee after a protected leave, and a job description that no longer reflects how the work is actually done may be of limited value as evidence.

Leave. Using an agent that has been trained on an employee's work to keep drawing on that output during the employee's leave may raise the question whether the employee is being required to work or should be compensated for their contribution. An agent's interim coverage of the employee's duties also does not by itself answer whether the employee must be reinstated at the conclusion of a leave.[142]

Employee-requested tools. Employees may ask to use their own AI tools as an accommodation. Such a request sits where the accommodation duty meets the confidentiality, data security, and trade secret risks of putting employer or third-party information into outside tools (see Sections IV.D and IV.I). If the employee's suggested tool cannot be used as requested, a question may arise whether an alternative, such as an enterprise-approved and controlled tool, would be effective.

3. Accommodating Employees in the Use of AI Itself

Employee-facing AI can create barriers of its own. Voice- or video-based tools, mandatory AI usage expectations, and activity monitoring may screen out or disadvantage employees with disabilities or pregnancy-related limitations, and those employees may request an alternative process or modified expectations as a reasonable accommodation, subject to undue hardship.[143] The applicant-side discussion in Section V.A raises a parallel issue: whether an individual receives notice of AI use early enough and is provided with a clear way to request an accommodation.[144] State and local enactments, such as California's automated-decision regulations, may also add requirements (see Section III.B).

Interactive process. An agent can collect information and schedule follow-ups, but the interactive process is a dialogue between the employer and the employee.[145] A question may arise as to whether an agent-run exchange satisfies that duty, or whether a human decision-maker must engage with the employee and own the outcome.

4. Records and Governance

Agent communications and decision logs may be the best evidence of when the employer received a request and how it responded. FMLA records must generally be kept for three years,[146] and the same records may be sought in discovery from employers and vendors. Issues to watch include how an agent handles a leave request embedded in a scheduling question, who is accountable for leave and accommodation decisions the agent supports, and when the agent is expected to hand off to HR. Leave, accommodation and reinstatement protections also vary by jurisdiction, so an agent deployed across states must comply with applicable rules in each.

I. Autonomy, Access, and Trade Secret Protection

An agentic AI tool's level of autonomy and level of an employer's systems and data amplifies how easily confidential information may leave a company and how quickly anyone notices. Trade secret law, which emerged around human actors and paper trails, still applies – although additional layers of protection may be necessary.

1. New Exits for Old Secrets

Companies have spent years closing the familiar exits for confidential information, such as forwarded emails, USB drives, and personal cloud accounts. AI opens new ones. An employee who pastes source code, a customer list, or a strategy deck into a public AI tool, meaning a free or consumer version not covered by a company contract, hands that data to an outside provider in seconds, and the provider's terms of use may let it keep or use the data. Some tools store what users enter, and users cannot always delete it.

Agents can act as insiders while also giving outsiders a way in. The Open Worldwide Application Security Project (OWASP), a nonprofit that publishes widely used security guidance, identifies the leading security risks for agents. These include tricking an agent with hidden instructions, misusing the programs an agent is connected to, an agent gaining more access than it was meant to have, data leaking out through what the agent does or says, corrupting the information an agent remembers from one task to the next, and giving an agent too much freedom to act without human checks.[147] The first of these, known as "prompt injection," remains the most common cause of agent security failures in actual use.[148] An instruction hidden in a document, calendar invite, or web page can carry as much weight with the agent as a command from its user. If an agent can reach private data, read content from outside the company, and send information out, a single planted instruction may lead it to send confidential data to an outsider without any employee doing anything wrong. OWASP's 2026 reporting describes incidents in which agents with poorly set access permissions and too much freedom to act sent data out of the organization.[149]

Information can leave without anyone copying a file. Requests to an AI tool that look harmless one at a time can, taken together, produce responses that rebuild a confidential product or process. AI also blurs the line between protected internal use and uncontrolled disclosure, and although it can generate valuable new information, it may also reduce the value of some existing secrets.

> "Confidentiality agreements, access controls, and training were largely built for a world of thumb drives, not instant uploads to AI platforms.

Autonomy complicates attribution. When an agent sits between the person who set it in motion and the place the information ends up, it may be unclear who acquired, disclosed, or used the information, and whether the conduct belongs to the individual employee or to the company that deployed the agent. Courts require plaintiffs to identify a trade secret with specificity and explain how the defendant improperly acquired, used, or disclosed it,[150] and in AI cases that burden can be difficult to meet. The problem is compounded by the fact that AI activity often leaves fewer records than email or file transfers, which makes the path of the information harder to reconstruct after the fact.

2. Same Law, Shifting Evidence

Every trade secret claim starts with identifying the secret. The Defend Trade Secrets Act (DTSA) protects a broad range of business, technical, and financial information if two conditions are met. The owner must have taken reasonable measures to keep it secret, and the information must derive independent economic value, actual or potential, from not being generally known or readily ascertainable through proper means.[151] The Uniform Trade Secrets Act and the state statutes modeled on it impose similar requirements.[152] In AI disputes, a company must say precisely what it is protecting and explain why secrecy gives that information value. The secret might be the data used to build an AI model, the settings the model learned during development, the instructions the company wrote to get useful results from AI tools, or ordinary business information that employees ran through those tools.

The reasonable measures requirement is where AI poses the most immediate threat. Information the owner did not reasonably protect is not a trade secret, however egregious the taking. The law does not demand absolute secrecy, only measures that are reasonable under the circumstances. Confidentiality agreements, access controls, and training were largely built for a world of thumb drives, not instant uploads to AI platforms. Entering trade secrets into a public AI tool may itself jeopardize their protection, and that what counts as reasonable must now account for the everyday work employees do with AI tools.[153]

Proving misappropriation is also harder when AI is involved. Under the DTSA, misappropriation includes acquiring a trade secret while knowing or having reason to know that it was acquired by improper means, as well as disclosing or using it without consent under circumstances the statute defines.[154] When an agent does the acquiring, that knowledge element can be difficult to tie to any particular person. Improper means include theft, bribery, misrepresentation, breach or inducement of a duty to maintain secrecy, and espionage through electronic or other means. Improper means do not include reverse engineering (taking apart a lawfully obtained product to learn how it works), independent derivation, or any other lawful means of acquisition.[155] That exclusion gives a defendant room to argue that a similar result came from public information run through an AI tool rather than from the plaintiff's secrets and assembling information lawfully available in the public domain generally does not create liability.

Misappropriation cases increasingly turn on AI evidence. Tracing where information went depends on what people typed into AI tools, what the tools produced, the records showing what actions an agent took, and activity in other programs the AI was linked to. Those records may show what a user submitted but not where the information went next, and activity in personal accounts or in AI tools used through a web browser may never reach the systems the company uses to keep records. Some records exist only briefly, or only if someone turned on record keeping in advance. An AI governance policy should address the ability to respond to incidents from the start. In practice, preservation letters and litigation holds should name AI accounts, the history of what users entered, and any programs linked to the AI, and computer forensic reviews should look for the records AI tools leave behind.

J. Vendors, Third Parties, and Accountability

A familiar principle of data privacy and technology governance is that an organization may outsource an activity, but not (necessarily) accountability for the result. Agentic AI gives that principle new salience. Depending on how a system is configured, an employer may be delegating not only the processing of employee information, but a measure of operational authority: the ability to pursue objectives, select methods, use tools, communicate with individuals, interact with other systems, and take actions without contemporaneous human direction or oversight.

1. Beyond Inaccurate Outputs

Traditional AI diligence often focuses on matters such as privacy, confidentiality, cybersecurity, reliability, bias, validation, and legal compliance. Those issues remain important. Agentic systems introduce an additional question: what might the system itself do in pursuit of its assigned objective?

An agent may perform satisfactorily in a demonstration or controlled evaluation without behaving identically in the customer's operating environment. Logs also remain critical, but should not automatically be treated as complete and independent evidence where the agent can access, influence, or circumvent the systems generating them, particular where agents act in groups and create language that is not in English or is difficult to follow. While we cannot assume that commercially deployed agents will ordinarily act deceptively or outside their authority, it is necessary to identify plausible failure modes that employers should consider before connecting agents to employment data, communications, or decision-making processes.

The practical question goes beyond whether an agent produces accurate outputs under expected conditions: it is whether the employer can constrain what the agent is authorized to do, observe what it actually does, interrupt it when necessary, and reconstruct its actions afterward.

2. Vendor Due Diligence

Employers have both a right and a responsibility to be demanding of agentic AI vendors. A vendor asking an employer to entrust its product with employee data, enterprise access, or authority within an employment process should expect enhanced scrutiny. If the vendor is surprised by that scrutiny, or cannot substantiate its assurances, that response is indicative of untrustworthiness.

Performing due diligence into an agent being considered for licensing and deployment should assess the exact product, version, configuration, integration, and proposed use. To determine whether an agentic product is "safe," "auditable," "explainable," or subject to "human oversight" as represented, employers can consider:

- the objectives and actions the agent is permitted and prohibited from pursuing;
- every system, tool, credential, data source, sub-agent, and material upstream component available to it;
- the actions requiring human approval and whether the reviewer has sufficient information and authority to intervene;
- how attempted and completed actions, tool calls (i.e., interactions with external tools, application programming interfaces (APIs), or systems), approvals, overrides, errors, and relevant versions are recorded;
- whether the agent can alter, suppress, or circumvent those records;
- how access can be restricted or withdrawn and the agent suspended immediately;
- how material model, configuration, tool, and safeguard changes will be disclosed and retested;
- how the vendor will support audits, individual explanations, correction requests, human review, investigations, litigation holds, and regulatory inquiries; and
- what remediation will be available if an action cannot be explained or reconstructed.

This need not be an adversarial exercise. Vendors developing sophisticated autonomous products should be prepared to work within operational constraints appropriate to the customer's environment.

3. The Current Regulatory Environment as a Floor for Compliance

Emerging AI laws provide important requirements but should not be treated as complete procurement specifications against which no more and no less action is taken. The technology, its capabilities, and the ways in which it is integrated into workplace systems may develop more quickly than detailed legal requirements.

A proposed deployment may also combine a model, objective, dataset, permissions, tools, integrations, and employment workflow that has never been tested as a whole. In that sense, the deployment may create an experimental operating environment even where the underlying product has been tested extensively. Material uncertainty is a reason for containment, instrumentation, graduated testing, and predefined stopping conditions, not a reason to assume that no further control is required because legislation is silent.

Multistate employers may therefore consider adopting a highest common denominator control framework nationally, supplemented by jurisdiction-specific legal requirements. This does not mean assuming that the most demanding state law applies in every jurisdiction, or that complying with the most burdensome law necessarily complies in every respect with other laws. It means selecting protective operational standards for consequential workplace uses instead of repeatedly designing each deployment around the minimum requirements of each location. The objective is to ensure that the authority entrusted to the agent does not exceed the ability to constrain, observe, interrupt, and account for its actions.

V. The Litigation and Investigation Environment

A. Trends in Employment-Related AI Litigation

Along with the rise in adoption of AI systems for use in various employment-related functions, lawsuits testing the use of AI against the backdrop of and employment- and privacy-related laws already in effect are also on the rise. The increase in use of AI tools in drafting litigation documents also presents challenges for litigants and courts.

Discrimination. The most prevalent theory of employment discrimination due to AI use proffered thus far by plaintiffs has been that AI use in hiring and screening causes disparate impact against protected groups by AI tools utilizing automated decision-making without sufficient human review. In the bellwether case of *Mobley v. Workday, Inc.*,[156] the plaintiffs allege that Workday's AI-powered job application screening tools discriminate against applicants on the basis of race, age, and disability in violation of federal and state anti-discrimination laws. *Mobley* is currently in the discovery phase and the issue of Workday's liability is still being adjudicated.

Disparate treatment theories of recovery are also in play. In 2022, the EEOC brought the very first lawsuit alleging disparate treatment on the basis of age in violation of the Age Discrimination in Employment Act through the use of AI tools.[157] The EEOC alleged that defendant iTutorGroup's hiring software automatically rejected female applicants aged 55 and older and male applicants aged 60 and older, affecting more than 200 U.S. applicants. While this matter settled in 2023 and it is therefore unknown how the EEOC's theory of recovery would have played out in court, this theory of liability will continue to be tested.

ADA and accessibility violations. Specific to disability discrimination concerns, the use of hiring technologies that require applicants to undergo algorithmically-based assessments, such as virtual interviews, personality tests, and other pre-hire examinations, implicate the Americans With Disabilities Act (ADA) and state-level accessibility laws. For example, if an employer utilizes facial and voice analysis technologies to evaluate applicants, otherwise qualified applicants with disabilities (*e.g.*, autism or speech impediments) may be unlawfully excluded from consideration.[158]

Additionally, the ADA and its state-level counterparts requires employers to provide reasonable accommodations to individuals with disabilities, including during the hiring process. This obligation may lead employers to take additional steps, such as giving prior notice to applicants regarding the use of AI hiring tools sufficient to allow applicants to decide to seek an accommodation and providing clear steps for applicants to do so without impacting their candidacy. As with other protected classes as discussed above, uniform use of an AI system that has the effect of screening out applicants with disabilities could give rise to a disparate impact suit.

Privacy issues. AI's reliance on knowledge bases necessarily relies on collection of data, which implicates the categories of data protected under consumer data privacy laws. For example, while AI note-taking tools can boost productivity and engagement, these tools introduce significant legal and operational risks, including potential violations of privacy and wiretap laws and exposure of confidential or privileged information.

Fair Credit Reporting Act. Another midcentury law, the federal Fair Credit Reporting Act (FCRA),[159] prohibits consumer reporting agencies from collection and dissemination of consumer information and providing the information to a third party without a specified use for it. In October 2024, the U.S. Consumer Financial Protection Bureau (CFPB) published guidance cautioning employers that their use of third-party consumer reporting companies to obtain background dossiers, algorithmic scores, and other third-party consumer reports about workers is subject to FCRA rules.[160] Thus, a plaintiff may allege an FCRA violation predicated on scraping social media websites, location data, internet activity, and other cookies and tracking data – without first providing the applicants the opportunity to opt out, without providing applicants a report of the information gathered, and without giving applicants the right to dispute the report results, as required under FCRA rules.

***Pro se* litigants and AI**. *Pro se* filings – those in which the plaintiff is unrepresented by counsel and proceeds on their own behalf – in federal employment cases nearly doubled from 2,052 cases in 2021 to 4,388 cases in 2025, according to Lex Machina's 2026 Employment Litigation Report. Pro se plaintiffs also made up 16.5% of all federal employment cases filed in 2025, up from only 9.7% in 2021.[161] While it cannot be determined with any specificity the degree to which AI played a role in influencing these figures, it is unquestionable that pro se plaintiffs commonly utilize AI tools in litigation, and are also the target audience for AI tools developed for that purpose.

While the outcomes of these pro se suits remain favorable to employers, with pro se plaintiffs losing on the merits at a ratio exceeding 40:1, they still trigger litigation burdens due to the voluminous filings these plaintiffs are able to generate through the use of legal AI tools. Skyrocketing litigation budgets in the defense of these cases is the primary concern for employers.

B. Theories of Liability Likely to Extend to Agentic Systems

Who bears responsibility when an agentic AI system makes consequential decisions and causes harm? Traditionally, liability has been tied to human conduct. The emergence of increasingly autonomous AI systems, however, challenges that framework and raises new legal and regulatory questions.

1. Enhanced Challenges of Assigning Liability with Agentic AI Systems

As discussed previously, agentic AI systems are typically distinguished by the following core characteristics: (1) autonomous planning, (2) real-world action through system integration, and (3) independent execution of actions. These features complicate the development of a liability framework for agentic AI because complex, multi-step reasoning can obscure how decisions are made; interactions with external platforms can introduce variables beyond the system's original design; and autonomous planning can produce outcomes that depart from intended parameters.

2. Developing Legal Framework

Although the law governing agentic AI is nascent, one principle is clear: using agentic AI does not, by itself, shield an employer from liability. As discussed previously, California and Connecticut law do not permit a defendant to avoid liability in an AI harm case by asserting that the AI system acted independently. The statutes preserve defenses based on causation, foreseeability, and comparative fault, suggesting that courts may allocate liability across the AI supply chain under established tort principles. In short, AI autonomy does not disrupt the chain of liability.

Employers outside of those states likewise should not assume that deploying agentic AI eliminates potential liability. No federal law excuses discrimination merely because a vendor, and not an HR manager, for example, made the decision in question. Under Title VII, employers must show that AI screening tools with a disparate impact are job-related and consistent with business necessity. The ADA imposes the same standard for tools that screen out individuals with disabilities, and also requires reasonable accommodations. The ADEA similarly covers tools that disproportionately exclude older workers. Under these laws, the employer remains liable for AI-assisted personnel decisions, without exception for how autonomously the AI system acted, including where it outsources substantial decision-making.

> Who bears responsibility when an agentic AI system makes consequential decisions and causes harm?

3. Who Can Be Held Liable and on What Grounds?

Companies that deploy agentic AI in their operations, products, or services may be liable for the agents' conduct. Potentially responsible parties often include the developer or vendor who designed the model; the business that deployed and configured it; and any third-party platform whose inadequate safeguards enabled the harm. Liability may begin with the deploying organization at the most fundamental level because it authorized the agent, set its permissions, and controlled its operating environment. Responsibility may, however, ultimately be partitioned among several parties, requiring courts to determine each party's respective contribution to the harm.

The primary legal theories of liability extending to agentic AI systems in an employment context span statutory discrimination, agency and vicarious liability, common law torts including negligence, and product liability.

Regulations. These potential liability theories operate against a broader regulatory backdrop. Although the United States lacks a comprehensive federal AI liability law, some states' laws regulate the deployment and use of AI tools in employment. In addition, federal anti-discrimination statutes are technology-neutral. Title VII, the ADA, and the ADEA apply to agentic AI systems just as they do to traditional screening methods.

Common law agency. Beyond these statutory requirements, traditional common law doctrines may also shape liability. Strictly speaking, an AI agent is not a legal agent because common law agency defines an agent as a "person." Plaintiffs may nevertheless argue for an analogous relationship. Although liability may be shared among developers, deployers, and others, employers cannot rely on AI autonomy as a defense, particularly in the states whose laws prevent this.

Vicarious liability. Related vicarious liability principles may provide a more direct basis for holding deployers responsible. When an AI agent acts for a company, courts may apply principles similar to those governing employer-employee relationships. If the agent acts within its authorized permissions, the deploying entity may be responsible, creating significant risk for companies that grant AI agents broad authority. Emerging case law has also answered the question by setting novel precedent that not only can legal principles like agency be applied to AI systems, but that AI *vendors/developers* (not just deployers) can also face direct liability for employment discrimination claims.[162]

Negligence. Where agency-based theories do not fully address the conduct, common law negligence may become an alternative basis for assigning liability for harm caused by AI agents. The familiar negligence framework applies: Did the deploying entity owe and breach a duty of care, and did that breach cause the harm? Liability may also extend to a third-party platform provider whose infrastructure enabled the harmful conduct by having inadequate safeguards in place.

Product liability. Product liability theories shift the focus from the deployer or platform provider to the developer by treating the agentic AI system as a potentially defective product. Because AI agents lack legal personhood, liability must ultimately rest with a person or organization.

C. Discovery Challenges Unique to Agentic Systems

When agentic AI influences employment decisions or workplace activities, litigation may turn on how the system acted, what information informed its actions, and who exercised meaningful oversight. Because these systems can carry out multistep tasks, interact with other systems, and operate with limited human intervention, decisions involving hiring, scheduling, performance management, or termination may be difficult to reconstruct.

Agentic AI may change the nature of discoverable information, but not the principles that govern discovery. Preservation, relevance, proportionality, and possession, custody, or control still apply. The open question is how those standards apply to electronically stored information (ESI) that agentic systems generate, modify, or discard as they operate. Courts are only beginning to address that question, and the answers will likely be fact-specific.

1. Explainability and Reconstructing System Activity

A challenged employment outcome may result from a sequence of system actions rather than a single output. Because agentic systems may rely on changing data and interactions across platforms, reconstructing what information was available and how the system reached an outcome can be difficult.[163] Reconstruction also involves people, and discovery may focus on who defined the system's objectives, approved its use, reviewed its outputs, and modified or overrode its recommendations.

System-generated explanations add another layer of complexity. Rationales, summaries, confidence scores, and reasoning logs may not fully reflect the underlying computational process. Some are contemporaneous records of system activity, while others are after-the-fact descriptions, and their evidentiary value will depend on how they were generated and whether they reliably reflect what the system actually did. Emerging explainability frameworks[164] may inform governance practices, but they do not by themselves establish legal standards or create discovery obligations.

Litigation may do more than regulation to open AI's "black box." Vendors and the employers that deploy their tools may have to explain under oath how those tools work.[165] Discovery may therefore become one of the primary means of explaining agentic AI decision-making.

2. Audit Trails and System Logs

Potentially relevant ESI includes prompts and instructions, permissions, retrieval queries, tool and API interactions, execution histories, agent-to-agent communications, human approvals or overrides, configuration and version records, and error logs. These records may help establish how a system operated and how humans interacted with it.

Logging, however, has practical limits. Logs may record actions without capturing the system's reasoning, the information available to it at the time, or the model, index, or configuration in use. Because logging capabilities often reflect vendor design choices rather than decisions by the organization using the tool, employers should understand what records their AI tools generate and retain. Organizations that follow AI governance practices may already maintain audit logs and approval records that become relevant in later disputes.

Where centralized or readable logs do not exist, that fact may inform discussions about proportionality and what can reasonably be reconstructed. And even where relevant data exists, it may not be reasonably accessible because of undue burden or cost. A party taking that position, however, must support it with specific facts rather than general claims about technical complexity, and a court may still order discovery if the requesting party shows good cause.[166]

Regulatory requirements may also generate relevant records, such as recordkeeping for certain high-risk employment uses under the EU AI Act, bias audits under New York City Local Law 144, and documentation under Colorado's AI statute. A record kept for regulatory purposes, however, is not automatically discoverable.

3. Preservation of Dynamic and Ephemeral Data

Agentic systems may create, modify, aggregate, overwrite, or delete information during ordinary operation, including session memory, rolling logs, retrieval caches, temporary files, and API records. The existing framework still governs. Under Federal Rule of Civil Procedure 37(e), the inquiry is whether relevant ESI that should have been preserved in anticipation of litigation was lost because a party failed to take reasonable steps, and whether it can be restored or replaced through additional discovery. Preservation requires reasonable, good-faith efforts, not every conceivable step, and responding parties are best situated to evaluate their own systems and preservation methods.[167]

In practice, employers should understand where AI-related data is stored, how long it is retained, and whether it can be preserved or exported, including from vendor-hosted systems. Preservation may need to reach beyond email and documents to prompts, logs, and other system data, with privacy, security, and trade secret concerns informing its scope. These are factors in assessing reasonableness, not fixed requirements, and what is reasonable will depend on the circumstances of each case. Cases involving auto-delete settings and transient data show the risk of not understanding these systems before data is lost.[168]

4. Possession, Custody, and Control

AI-related records may be distributed between employers and technology providers. Rule 34's control analysis varies by jurisdiction and may consider legal rights or practical ability to obtain records.[169] Vendor agreements addressing access, retention, export, and legal holds therefore warrant attention, although contractual terms alone do not determine control. Recent preservation orders directing AI providers to retain user output logs show that vendor-held data is increasingly a focus of discovery.[170]

5. Protective Orders and Confidential Information

When agentic systems become the focus of discovery, protective orders may need provisions beyond standard confidentiality terms. Source code, model configurations, system prompts, orchestration logic, and other vendor technical information may be trade secrets, and parties may need to decide whether that material calls for heightened designations, such as attorneys' eyes only or source code tiers, or controlled inspection protocols that limit copying and access. Protective orders may address how parties and their vendors may use AI tools with produced materials. They could prohibit uploading confidential, privileged, or sensitive employee information to tools that lack adequate safeguards against retention, disclosure, or unauthorized training. These issues may need to be addressed early, particularly when outside litigation vendors or third-party AI platforms are involved.[171]

D. How Plaintiffs May Adapt Existing Employment Claims to Agentic AI

The legal theories may be familiar, but agentic AI could change how plaintiffs build cases. An agent may gather information, apply criteria, communicate with employees, and initiate actions with limited human involvement. Before filing suit, plaintiffs' counsel may examine how the technology was used, how the employee interacted with it, where a human entered the process, and what records the system created.

1. From a Single Decision to a Decision Chain

Cases involving automated hiring tools already show how existing legal theories can apply when technology plays a role in an employment decision. In *EEOC v. iTutorGroup, Inc.*, the EEOC alleged that recruiting software automatically rejected applicants at specified age thresholds.[172] New York City's automated-employment-decision-tool (AEDT) law similarly treats screening as part of the employment decision, even if the tool does not make the final hiring or promotion decision.[173]

Agentic AI may complicate that analysis because the challenged employment action may be the product of several decisions made along the way, rather than a single automated decision. Depending on the claim, any discrete decisions, or the process as a whole, may become relevant to the plaintiff's theory.

2. Reframing the Decisionmaker and Human Review

Agentic AI may also make it harder to identify who, or what, made the challenged decision. A plaintiff may look beyond the manager who communicated the final outcome and examine the roles of the people who selected or configured the system, the individuals responsible for reviewing its work, and the vendor that developed or operated it. Plaintiffs may

use familiar agency principles to argue that responsibility extends beyond the employer that ultimately makes or communicates the decision.

California has begun to address responsibility and human review directly. Civil Code section 1714.46 provides that a person who develops, modifies, or uses AI cannot avoid a civil action simply by arguing that the AI acted autonomously, although traditional defenses involving causation, foreseeability, and comparative fault remain available.[174] In addition, the No Robo Bosses Act will prohibit employers from relying solely on an automated decision system to discipline or terminate an employee and will require human corroboration when the employer relies primarily on such a system.[175]

These laws also illustrate a broader question plaintiffs are likely to ask: how much independent judgment did the human reviewer actually exercise? Counsel may examine what information the reviewer received, whether the reviewer could evaluate the basis for the recommendation, whether the reviewer had authority to reach a different result, and whether that authority was exercised. Approval and override records may show how the review process operated in practice, rather than only how it was designed.

3. Following the Agent Back to the Data

Plaintiffs may look beyond the decision to the information the agent collected, combined, or inferred. Existing laws may provide different ways to challenge that information or its use. In the hiring context, for example, plaintiffs are already testing whether existing laws governing candidate selection, such as state and federal antidiscrimination laws and consumer background check laws can reach AI-generated assessments. The theory will depend on the underlying law, but the factual inquiry is similar: what did the system know, how did it use that information, and did that use affect what happened next?

Agentic systems may also leave behind records that allow those questions to be explored in greater detail. Depending on how a system operates, potentially relevant evidence may include the instructions given to the agent, the sources of information it accessed, scores or rankings it generated, communications it sent, changes to its configuration, and records showing when a human approved, rejected, or changed its recommendation. Emerging laws that require employers to specify when a layoff or other restructuring is due to automation may make some of that information visible even before litigation begins. For example, California law will require Cal-WARN notices involving workforce reductions caused in whole or substantial part by AI or other automated technology to include additional information about the work being affected, the functions being automated, and the technology involved.[176] Those disclosures do not establish that an employer acted unlawfully, but they may help plaintiffs identify the role technology played and shape the records they seek if litigation follows.

Similar issues may arise in a unionized workplace. If an agent is used to schedule employees, monitor performance, or provide information used in discipline, its implementation or use may implicate existing bargaining obligations depending on the circumstances and the applicable collective bargaining agreement. An agent may also generate information that becomes relevant to a union information request or to a later challenge to discipline. The focus will be on what function the agent performed and whether that function affected employees' wages, hours, or other terms and conditions of employment.

In this new era, plaintiffs may move beyond the reasons given by the manager who made or communicated the final decision and examine the automated and human decisions that preceded it, who controlled those decisions, and what information the system used. As agents take on more of those intermediate steps, the litigation may increasingly turn on how the decision was reached, not simply the final decision itself.

VI. Key Takeaways

As discussed throughout this Report, agentic AI raises familiar employment law questions in new forms. Four themes run through them:

Governance and use policies.

Many employers can extend existing AI governance structures to agents instead of starting over. For agents, policies may add what actions an agent is authorized to take, which systems and data it may access, when it must hand off actions to a person, and how its activity is logged. They may also add a central inventory of agents and a named accountable owner with authority to slow or stop a deployment.

Compliance themes.

AI-specific employment laws differ in scope and detail, but several obligations recur across jurisdictions: notice, human review and appeal, bias testing, recordkeeping, explanation of adverse actions, and access to and correction of data. Laws not written for AI, including federal anti-discrimination statutes, apply whether or not AI is involved. Mapping each agentic workflow against both the AI-specific rules and the general employment statutes may help compliance teams see where obligations overlap.

AI literacy.

AI literacy is becoming both a regulatory expectation and a practical safeguard. The EU AI Act, as replaced by the Digital Omnibus on AI, requires providers and deployers to take measures to support the AI literacy of their staff. Under California's ADMT regulations, human review keeps a decision outside the rules only if the reviewer can interpret the output, considers other relevant information, and has authority to change the decision. Training for agentic systems may cover what each agent is authorized to do, how to recognize and escalate errors, and when to override an agent.

Vendor management.

Diligence may begin with what a product actually does. Although the law on vendor liability is developing, it does not displace the employer's own obligations. Key terms may need updating for agentic tools, potentially including compliance representations, cooperation with risk assessments and regulatory inquiries, limits on data use and retention, and audit rights; notice when the vendor updates or retrains a model, independent audit rights over training data and outputs, and indemnification that addresses regulatory fines and class exposure, given the authors' view that standard AI vendor indemnities "are often functionally illusory"; and cooperation in preserving logs, notice of incidents, limits on the vendor's use of company data for training, and review of the vendor's own insurance.

In addition to these core issues, emerging issues to monitor include:

> **Agent accountability frameworks.** As agents take on more tasks, a recurring question is who answers for what they do. Current sources point in one direction. California has codified one piece of this: since January 1, 2026, a defendant that developed, modified, or used AI may not assert as a defense that the AI autonomously caused the harm. Standards work is also under way, including NIST's AI Agent Standards Initiative, launched in February 2026, which covers agent identity and authorization. Section IV.B addresses agency doctrine in detail.

> **Cross-border and multijurisdictional compliance.** Agentic HR systems are often deployed centrally, so a single agent may fall under several regimes at once. For multinational employers, a single inventory of where each agent operates and whose data it touches may be a practical starting point. Section II.C addresses these regimes in detail.

> **Insurance and indemnification.** Insurance markets are adjusting. Insurance Services Office (ISO) endorsements effective January 1, 2026, allow commercial general liability insurers to exclude claims arising out of generative AI (CG 40 47, CG 40 48, and CG 35 08), at least one carrier has introduced an "absolute" AI exclusion for directors and offices (D&O), errors and omissions (E&O), and fiduciary liability policies,[177] and some insurers use broad AI exclusions in employment practices liability policies. According to a September 2026 CSIS analysis, ISO confirmed in July 2026 that it is drafting exclusions specific to agentic AI.[178] Standalone AI liability coverage exists but remains limited, and standard AI vendor indemnities may be capped at contract value or carve out key risks. Employers may therefore want to review how AI-related claims would map to their EPLI, D&O, cyber, and E&O coverage, and to their vendor indemnities, well before renewal.

A. A Proactive Approach to Agentic AI: Decide, Design, Defend

Agentic capabilities can change familiar functions and workflows. Employers can make their organizations more efficient, effective, compliant, and defensible through AI-enabled workforce transformation. A proactive approach, with a foundation in the legal considerations addressed in this Report, can increase the value and impact of technology investments. The following components can be adapted to the organization, risk tolerance, jurisdictions, and workflows.



1. Decide: Identify Opportunities and Assess Readiness

Agentic transformation may start with an inventory of candidate workflows and a clear view of the organization's readiness to govern them. For each workflow, two questions can guide the decision:

- What level of autonomy would the agent have using the spectrum described in Section I.B: recommend, draft, execute with approval, or execute autonomously?
- What risk factors are present?

The answers support a measured approach to increasing agent autonomy based on organizational capability and governance

2. Design: Build Compliant and Effective Systems

Once a use case and autonomy level are chosen, design choices shape both compliance and value. Seven elements of legal significance recur across the sources reviewed:

- **Governance framework**, including ownership, approval gates, permissions, and escalation paths for each agent.
- **Human oversight checkpoints**, especially where it can change outcomes, with appropriate documentation.
- **Vendor due diligence**, with particular attention to model-change notice, audit rights, access to testing data, and log preservation.
- **Bias auditing strategy**, reflecting quality, efficacy, recency, scope, results, the response, and privilege protections.
- **Documentation and audit trails**, accounting for legal obligations, anticipated uses, and retention costs and administration.
- **Training** designed for users to understand what each agent does and when to intervene, including role-based training with the AI literacy efforts.

3. Defend: Prepare for Scrutiny

Defending agentic decisions begins well before any claim. Three proactive steps can support defense efforts:

- Prepare for AI prompts, outputs, and agent logs to potentially be discoverable, and update litigation hold templates and custodian interviews accordingly.
- Consider the utility of contemporaneous records of the pre-deployment risk analysis, the allocation of responsibilities, and the monitoring plan to demonstrate good faith and reasonable efforts to avoid harm.
- Determine an approach for AI-use registers, potentially listing tools, owners, vendors, data inputs, and decisions influenced.

Conclusion

Agentic AI gives employers an opportunity to reimagine how work gets done. Approached strategically and grounded in the considerations this Report addresses, AI-enabled workforce transformation can deliver durable adoption and organizational value: decide deliberately, design carefully, and be ready to defend both.

Endnotes

1 Gartner, *2026 Hype Cycle for Agentic AI* (Apr. 15, 2026), reporting results of the 2026 Gartner CIO and Technology Executive Survey.

2 Press Release, Gartner, Inc., Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027 (June 25, 2025).

3 Deloitte AI Inst., *State of AI in the Enterprise: The Untapped Edge* 17, 20 (Jan. 2026).

4 Press Release, Gartner, Inc., Gartner Identifies Six Steps to Manage AI Agent Sprawl (Apr. 28, 2026).

5 Gartner, *supra* note 1.

6 National Institute of Standards and Technology, *Agentic AI* (https://www.nist.gov/agentic-si).

7 International Association of Privacy Professionals, *AI governance in the agentic era* (July 16, 2025).

8 U.S. Government Accountability Office, *Science & Tech Spotlight: AI Agents*, GAO-25-108519 (September 2025).

9 *Id.*

10 IBM, *Generative AI vs. Predictive AI: What's the Difference?* (August 9, 2024); National Institute of Standards and Technology, *Agentic AI* (https://www.nist.gov/agentic-si).

11 National Institute of Standards and Technology, *Computer Security Resource Center Glossary* (https://csrc.nist.gov/glossary).

12 National Institute of Standards and Technology, *AI Use Taxonomy: A Human-Centered Approach*, NIST AI 200-1 (March 2024).

13 National Institute of Standards and Technology, *Computer Security Resource Center Glossary* (https://csrc.nist.gov/glossary).

14 U.S. Government Accountability Office, *Science & Tech Spotlight: Generative AI*, GAO-23-106782 (June 13, 2023).

15 National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1 (July 2024).

16 National Institute of Standards and Technology, *The Language of Trustworthy AI: An In-Depth Glossary of Terms*, NIST AI 100-3 (March 2023).

17 IBM, *What Is Computer Vision?* (September 17, 2025).

18 U.S. Government Accountability Office, *Digital Surveillance of Workers: Tools, Uses, and Stakeholder Perspectives*, GAO-24-107639 (August 28, 2024)

19 See Cloud Sec. Alliance AI Safety Initiative, *Agentic AI Autonomy Levels and Control Framework* 9 (ver. 2.0, 2026).

20 This report primarily uses the term "decision-making," but "decisionmaking" and "decision making" may also appear where those terms are part of a statutory or regulatory definition.

21 See Colo. Rev. Stat. § 6-1-1701(15); Regulation 2024/1689, of the European Parliament and of the Council of 13 June 2024, Laying Down Harmonised Rules on Artificial Intelligence, art. 14(4), 2024 O.J. (L 1689) 1 (EU) (requiring that persons assigned to human oversight be able to interpret the system's output, remain aware of "automation bias," and override or reverse the output).

22 Cloud Sec. Alliance AI Safety Initiative, *supra* note 1, at 9–13.

23 N.Y.C. Admin. Code § 20-870.

24 *E.g.,* Cal. Code Regs. tit. 11, § 7001(e)(1)(A)–(C).

25 Korn Ferry, *Talent Acquisition Trends 2026* (Oct. 28, 2025).

26 Gartner, press release, *Gartner Survey Finds 91% of Customer Service Leaders Under Pressure to Implement AI in 2026* (Feb. 18, 2026).

27 McKinsey & Company, *The State of AI: Global Survey* (Aug. 25, 2026).

28 Press Release, Gartner, Gartner Predicts Agentic AI Will Autonomously Resolve 80% of Common Customer Service Issues Without Human Intervention by 2029 (Mar. 5, 2025).

29 Press Release, EY, EY Survey Reveals That Technology Companies Are Setting the Pace of Agentic AI (May 2025).

30 National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework* (AI RMF 1.0) (2023), GOVERN 3.2 and MANAGE 4.1; NIST, *Generative Artificial Intelligence Profile*, NIST AI 600-1 (2024).

31 Gallup, *AI in the Workplace: What Separates Adopters and Holdouts* (Apr. 12, 2026).

32 *Id*.

33 N.Y.C. Admin. Code § 20-870; N.Y.C. Department of Consumer and Worker Protection, Automated Employment Decision Tools Frequently Asked Questions; Regulation (EU) 2016/679, art. 22.

34 Deloitte, *The State of AI in the Enterprise* (2026).

35 Press Release, Gartner, Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027 (June 25, 2025).

36 EEOC, Employment Discrimination and AI for Workers

37 Jim Paretti, Chris Gokturk, and Alyesha Asghar, *Trump Administration Moves to Eliminate Federal Government's Use of Disparate Impact Theory Liability*, Littler ASAP (April 24, 2025).

38 Jim Paretti, *EEOC to Close Investigations of Disparate Impact Discrimination*, Littler ASAP (Sept. 25, 2025).

39 Executive Order 14179, Removing Barriers to American Leadership in Artificial Intelligence (Jan. 23, 2025).

40 *America's AI Action Plan*, https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf

41 Press release, U.S. Department of Labor, *US Department of Labor applauds President Trump's 'AI Action Plan' to achieve global dominance in artificial intelligence* (July 23, 2025).

42 *Id*.

43 Executive Order 14365, *Ensuring a National Policy Framework for Artificial Intelligence* (Dec. 11, 2025).

44 *See, e.g., New York Central R.R. Co. v. Chisholm*, 268 U.S. 29, 31 (1925); *North Alaska Salmon Co. v. Pillsbury*, 174 Cal.1, 4 (1916).

45 *Edgar v. MITE Corp*., 457 U.S. 624, 642-43 (1989).

46 *See Healy v. Beer Institute, Inc*., 491 U.S. 324, 336–37 (1989); *Pike v. Bruce Church, Inc*., 397 U.S. 137, 142 (1970); *Kearney v. Salomon Smith Barney, Inc*., 39 Cal.4th 95 (Cal. 2006).

47 *See Flynt v. Bonta*, 131 F.4th 918 (9th Cir. 2025); *Edgar*, 457 U.S. 624 at 643-44.

48 29. C.F.R. Part 1607.

49 820 Ill. Comp. Stat. 42/20.

50 775 Ill. Comp. Stat. 5/2-101, 5/2-102.

51 Colo. Rev. Stat. §§ 6-1-1701 - 6-1-1709.

52 Conn. P.A. 26-15, §§ 9, 10.

53 Conn. Gen. Stat. § 46a-60.

54 Cal. Code Regs. tit. 2, §§ 11008 *et seq*.

55 2026 Cal. Stat. ch. 859 (S.B. 947) (adding Cal. Lab. Code §§ 1520–1526.7).

56 SB 947 also requires a stand-alone post use notice when an automated decision system is used in a disciplinary or termination decision.

57 775 Ill. Comp. Stat. 5/2-101, 5/2-102.

58 *See* Cal. Code Regs. tit. 2, §§ 11008 *et seq*; 2026 Cal. Stat. ch. 859 (S.B. 947) (adding Cal. Lab. Code §§ 1520–1526.7).

59 Conn. Gen. Stat. § 46a-60.

60 *See* Colo. Rev. Stat. §§ 6-1-1701 - 6-1-1709.

61 *See* Thomas B. Nachbar, *Algorithmic Fairness, Algorithmic Discrimination*, 48 Fla. St. U. L. Rev. 509, 530 (2021).

62 42 U.S.C. § 2000e-2(a).

63 Keith E. Sonderling, Bradford J. Kelley & Lance Casimir, *The Promise and The Peril: Artificial Intelligence and Employment Discrimination*, 77 U. Mia. L. Rev. 1, 23 (2022).

64 Alex Engler, *Auditing Employment Algorithms for Discrimination*, Brookings Institution (March 12, 2021).

65 Kyra Wilson & Aylin Caliskan, *Gender, Race, and Intersectional Bias in AI Resume Screening via Language Model Retrieval*, Brookings Institution (April 25, 2025).

66 *Id.*

67 42 U.S.C. § 2000e-2(k).

68 Sonderling et al., *supra*, at 24 (citing Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 Wm. & Mary L. Rev. 857, 874-75 (2017).

69 *See* Ayelet Israeli & Eva Ascarza, *When Bias Creeps into AI, Managers Can Stop It by Asking the Right Questions*, Harvard Business School Working Knowledge (October 18, 2022).

70 Federal Trade Commission, *Combatting Online Harms Through Innovation: A Report to Congress* (June 2022).

71 *See, e.g.,* the Illinois Biometric Information Privacy Act, 740 Ill. Comp. Stat 14/5 *et seq.*

72 Md. Code, Lab. & Emp. § 1202; Portland, Or. City Code §§ 34.10.010 – 34.10.050.

73 *See* 18 U.S.C. §§ 1836, 1839. U.S. Patent and Trademark Office, *Trade Secret Policy* (July 29, 2026).

74 *See* 15 U.S.C. § 1125(a).

75 National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework 1.0*, NIST AI 100-1 (January 2023).

76 National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1 (July 2024).

77 42 U.S.C. §§ 2000e(b), 12111(5)(A); 29 U.S.C. § 630(b).

78 Anthropic, Building Effective AI Agents (https://www.anthropic.com/research/building-effective-agents).

79 EEOC Compliance Manual § 2-III(B)(2)(b).

80 Restatement (Third) of Agency § 7.03.

81 Deborah DeMott, *When Agentic AI Met the Common Law of Agency,* Duke Law School Public Law & Legal Theory Series No. 2026-35, (May 27, 2026)..

82 Deven R. Desai and Mark Riedl, *Using Agency Law to Tame AI Agents*, Berkeley Technology Law Journal, (July 30, 2025).

83 *General Building Contractors Ass'n, Inc. v. Pennsylvania* 458 U.S. 375 (1982).

84 Cal. Code Regs. tit. 2, § 11008(b).

85 *Id.* at §§ 11008(b), 11009(b).

86 *Fiano v. Old Saybrook Fire Co. No. 1, Inc.*, 332 Conn. 93, 102–03 (2019) (internal quotation marks omitted).

87 *Bentley v. AutoZoners, LLC*, 935 F.3d 76, 91 (2d Cir. 2019) (*quoting Vance v. Ball State University*, 570 U.S. 421, 431 (2013).

88 Equal Emp. Opportunity Comm'n, *Enforcement Guidance: Vicarious Employer Liability for Unlawful Harassment by Supervisors* § IV.A (June 18, 1999); *See also*, *Burlington Industries, Inc. v. Ellerth*, 524 U.S. 742, 762-65 (1998); *Faragher v. City of Boca Raton*, 524 U.S. 775, 807-08 (1998).

89 N.Y.C. Admin. Code § 8-107(13).

90 See Ivie A. Serioux & Jerry Zhang, *Digital Doppelgängers in the Workplace: Emerging Employment Law Risks*, N.Y. State Bar Ass'n (Sept. 16, 2026) (discussing whether AI systems operating with apparent supervisory authority may be analyzed under existing employer-liability frameworks).

91 NIST, *NIST AI Risk Management Framework 1.0, Appendix C: AI Risk Management and Human-AI Interaction* (2023).

92 42 U.S.C. §2000e-2(k).

93 See, e.g.,A. J. Wu, R. Liu, X. Bai, and T.L. Griffin, *Large Language Models Develop Novel Social Biases Through Adaptive Exploration*, ICML (Nov. 8, 2026).

94 42 U.S.C. 2000e-2(k)(1)(A).

95 740 Ill. Comp. Stat. 14 *et seq.*

96 Cal. Civ. Code § 1798.100 *et seq.*

97 *See, e.g.,* Cal. Bus. & Prof. Code §§ 22625 *et seq.*

98 2026 WL 2351237 (N.D. Cal. 2026).

99 See 28 C.F.R. Part 202.

100 Department of Justice, *Data Security Program: Compliance Guide* (Apr. 11, 2025).

101 The Department of Labor briefly addressed workplace AI in Field Assistance Bulletin No. 2024-1, but later withdrew the bulletin after the executive order 14110 "Safe, Secure and Trustworthy Development and Use of Artificial Intelligence" under which it was issued was revoked January 20, 2025 by President Trump. The bulletin therefore is not current DOL guidance. Nevertheless, the statutory and regulatory wage and hour obligations implicated by automated scheduling and task allocation remain in place.

102 29 U.S.C. § 157.

103 *See, e.g., NCRNC, LLC v. NLRB*, 94 F.4th 67 (D.C. Cir. 2024).

104 *Oakwood Healthcare, Inc.,* 348 N.L.R.B. 686, 689–93 (2006) (interpreting the terms "assign," "responsibly to direct," and "independent judgment" under Section 2(11)).

105 *NLRB v. Katz*, 369 U.S. 736, 743 (1962); 29 U.S.C. §§ 158(a)(5), 158(d).

106 *See First Nat'l Maint. Corp. v. NLRB*, 452 U.S. 666 (1981).

107 *Colgate-Palmolive Co.*, 323 NLRB 515 (1997); *see also Brewers & Maltsters, Local Union No. 6 v. NLRB*, 414 F.3d 36 (D.C. Cir. 2005).

108 Cal. Exec. Order N-6-26 (2026); Joy Rosenquist, Alex MacDonald, Britney Torres & Alice Wang, *California Executive Order Focuses on Worker Protections for Disruption Caused by AI*, Littler ASAP (May 29, 2026).

109 29 U.S.C. §§ 2101(a)(1)–(3), 2102(a).

110 29 U.S.C. § 2102(d).

111 2026 Cal. Stat. ch. 860 (S.B. 951); *see* Cal. Lab. Code § 1401(d)(1). California's Cal-WARN statute is codified at Cal. Lab. Code § 1400 *et seq.*

112 *See Reeves v. Sanderson Plumbing Prods., Inc.*, 530 U.S. 133, 147–48 (2000).

113 *Fibreboard Paper Prods. Corp. v. NLRB*, 379 U.S. 203, 215 (1964); 29 U.S.C. § 158(a)(5), (d).

114 *First Nat'l Maint. Corp. v. NLRB*, 452 U.S. 666, 679 (1981).

115 42 U.S.C. § 2000e-2(k); *Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971).

116 *Smith v. City of Jackson*, 544 U.S. 228, 240 (2005); *Meacham v. Knolls Atomic Power Lab.*, 554 U.S. 84, 87 (2008) (a reduction-in-force case); 29 U.S.C. § 623(f)(1).

117 29 C.F.R. § 1625.7(e).

118 Cal. Code Regs. tit. 2, § 11008 *et seq.*; Colo. Rev. Stat. §§ 6-1-1701 to -1709.

119 Cal. Code Regs. tit. 2, § 11017(e) [cite-check: subsection taken from a pleading; confirm against the current regulation].

120 29 C.F.R. § 825.220(c); *Bachelder v. Am. W. Airlines, Inc.*, 259 F.3d 1112, 1124–25 (9th Cir. 2001).

121 *See, e.g.*, Cal. Gov't Code § 12945.2 (California Family Rights Act); Wash. Rev. Code § 50A.40.010.

122 42 U.S.C. §§ 12112(b)(5), 2000e(k), 2000gg-1.

123 29 C.F.R. § 825.216(a).

124 29 C.F.R. § 825.220(b).

125 See 42 U.S.C. § 12112(d)(3)(B), (d)(4)(C).

126 29 U.S.C. § 626(f)(1)(F)(ii), (H); 29 C.F.R. § 1625.22(f).

127 See 29 C.F.R. § 1625.22(a)(3), (b)(4)–(5).

128 See *Oubre v. Entergy Operations, Inc.*, 522 U.S. 422, 426–28 (1998).

129 See *Mobley v. Workday, Inc.*, 740 F. Supp. 3d 796 (N.D. Cal. 2024).

130 Cal. Civ. Code § 1714.46(b).

131 See 29 C.F.R. § 1607.4(D) (Uniform Guidelines on Employee Selection Procedures).

132 29 U.S.C. § 2601 *et seq.* (FMLA); 42 U.S.C. § 12101 *et seq.* (ADA); 42 U.S.C. § 2000gg *et seq.* (PWFA); Cal. Gov't Code § 12945.2 (CFRA); Cal. Gov't Code § 12940 (FEHA).

133 29 C.F.R. §§ 825.302(c), 825.303(b) (FMLA); 29 C.F.R. § 1630.2(o)(3) (ADA interactive process).

134 29 C.F.R. § 825.300(b)(1).

135 29 C.F.R. § 825.220(b); *see* 29 U.S.C. § 2611(2) (eligible employee).

136 29 C.F.R. § 825.220(c).

137 29 C.F.R. § 825.214.

138 42 U.S.C. § 12112(d)(4)(C); 29 C.F.R. § 825.500(g).

139 See 42 U.S.C. § 12112(b)(6); 42 U.S.C. § 2000gg-1; 29 C.F.R. pt. 1636.

140 42 U.S.C. §§ 12111(10), 12112(b)(5)(A); 29 C.F.R. pt. 1630, app. § 1630.9.

141 42 U.S.C. § 12111(8).

142 29 C.F.R. §§ 825.214, 825.216(a).

143 42 U.S.C. § 12112(b)(5)(A), (b)(6); 42 U.S.C. § 2000gg-1.

144 See Bradford Kelley & Sean O'Brien, *The ADA Turns 34: The Intersection of Technology, AI, and Individuals with Disabilities*, Littler ASAP (Aug. 21, 2024).

145 29 C.F.R. § 1630.2(o)(3).

146 29 C.F.R. § 825.500(b).

147 OWASP GenAI Security Project, *OWASP Top 10 for Agentic Applications for 2026* (Dec. 9, 2025).

148 *Id.*

149 *Id.*

150 *See, e.g., Diodes, Inc. v. Franzen*, 260 Cal. App. 2d 244, 252, 67 Cal. Rptr. 19 (Cal. Ct. App. 1968).

151 18 U.S.C. § 1839(3)(A)–(B).

152 Unif. Trade Secrets Act § 1(4)(i)–(ii) (Unif. L. Comm'n 1985); *see, e.g.*, Cal. Civ. Code § 3426.1(d)(1)–(2).

153 Camilla A. Hrdy, *Trade Secrecy Meets Generative AI*, 100 Chi.-Kent L. Rev. 317 (2025).

154 18 U.S.C. § 1839(5)(A)–(B).

155 *Id.*

156 Case No. 2:23-cv-007700, filed February 21, 2023.

157 *EEOC v. iTutorGroup, Inc. et al.*, Case No. 1:22-cv-02565 (E.D.N.Y.).

158 See Bradford Kelley and Sean O'Brien, *The ADA Turns 34: The Intersection of Technology, AI, and Individuals with Disabilities*, Littler ASAP (August 21, 2024).

159 15 U.S.C. § 1681 *et seq.*

160 Consumer Financial Protection Bureau, *Consumer Financial Protection Circular 2024-06 Background Dossiers and Algorithmic Scores for Hiring, Promotion, and Other Employment Decisions* (October 24, 2024).

161 LexisNexis, *Labor and Employment Federal Litigation Trends 2026* (March 10, 2026).

162 *See Mobley v. Workday, Inc.*, 740 F. Supp. 3d 796 (N.D. Cal. 2024).

163 *See* Tara S. Emory & Maura R. Grossman, *The Next Generation of AI: Here Come the Agents!*, Nat'l L. Rev. (Dec. 30, 2024).

164 See NIST, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (Jan. 2023); Yonadav Shavit *et al.*, *Practices for Governing Agentic AI Systems* (OpenAI, Dec. 2023).

165 *See, e.g., Mobley v. Workday, Inc.*, 740 F. Supp. 3d 796 (N.D. Cal. 2024).

166 *See* Fed. R. Civ. P. 26(b)(2)(B).

167 *See* The Sedona Principles, Third Edition, 19 Sedona Conf. J. 1 (2018), Principles 5 and 6.

168 *See, e.g.*, *DR Distributors, LLC v. 21st Century Smoking, Inc.*, 513 F. Supp. 3d 839 (N.D. Ill. 2021); *Fast v. GoDaddy.com LLC*, 340 F.R.D. 326 (D. Ariz. 2022).

169 See The Sedona Conference, *Commentary on Rule 34 and Rule 45 "Possession, Custody, or Control,"* 17 Sedona Conf. J. 467 (2016).

170 *In re OpenAI, Inc., Copyright Infringement Litigation*, No. 23-CV-11195, 2025 WL 1443780 (S.D.N.Y. May 16, 2025).

171 *EEOC v. iTutorGroup, Inc. et al.*, Case No. 1:22-cv-02565 (E.D.N.Y.).

172 *EEOC v. iTutorGroup, Inc. et al.*, Case No. 1:22-cv-02565 (E.D.N.Y.).

173 N.Y.C. Admin. Code § 20-870 (2026); N.Y.C. Dep't of Consumer & Worker Prot., *Automated Employment Decision Tools: Frequently Asked Questions 1–2* (June 29, 2023).

174 Cal. Civ. Code § 1714.46.

175 2026 Cal. Stat. ch. 859 (S.B. 947) (adding Cal. Lab. Code §§ 1520–1526.7).

176 2026 Cal. Stat. ch. 860 (S.B. 951).

177 Anthony Crawford, *How Policyholders Can Press for Narrower AI Exclusions*, Law360 (Aug. 7, 2026).

178 Gregory C. Allen, *The Insurance Industry's Retreat from AI Threatens to Slow Innovation and Adoption, Ctr. for Strategic & Int'l Stud.* (Sept. 4, 2026).

Appendix A.
Compendium of U.S. State Laws and International Laws Governing Use of AI in the Workplace

I. State and Local Laws in the U.S.

The laws already passed or in effect in the employment space in the United States at the state and local level tend to fall into one of a handful of categories: regulation of AI use in employment decisionmaking, limitations on AI "bossware" monitoring tools, business restructuring requirements when workers are displaced due to technology, and the parameters of liability for AI-caused harms.

A. Laws Regulating Automated Decisionmaking in Employment

California. In September 2026, the state enacted a package of legislation that regulates how employers may use AI. Senate Bill 947, which takes effect July 1, 2027, creates new obligations under the California Labor Code and prohibits an employer from using an automated decision system (ADS) to do any of the following: (1) prevent compliance with or violate any federal, state, or local labor, occupational health and safety, employment, or civil rights laws or regulations; (2) infer an employee's protected status under the California Fair Employment and Housing Act; or (3) predict and take adverse action against a worker for exercising their legal rights, including but not limited to rights guaranteed by state and federal employment and labor law. Further, the law expressly prohibits an employer from relying solely on an ADS when making a disciplinary or termination decision. If an employer primarily relies upon an ADS output to make a disciplinary or termination decision, the law requires a human in the loop. The employer must direct a human to corroborate the decision using the data that was collected or used to produce the ADS output or other relevant corroborating or supporting information. If an employer cannot corroborate the ADS output, or the human reviewer has concluded that the ADS output is inaccurate, incomplete, or misleading, the employer cannot use the ADS output to make a disciplinary or termination decision.

In addition to these statutory requirements under the Labor Code, the administrative regulations implementing two state statutes, the Fair Employment and Housing Act (FEHA) and the California Consumer Privacy Act (CCPA), contain significant compliance obligations regulating the use of automated decisionmaking systems. The FEHA is a traditional antidiscrimination statute – it prohibits discrimination in employment on the basis of an applicant's or employee's protected characteristics – and does not itself reference AI. In 2025, the California Civil Rights Council amended the FEHA regulations to address how the FEHA applies in the case of algorithmic discrimination.

Broadly speaking, the regulations establish that it is unlawful for an employer or other covered entity to use an automated decision system or selection criteria (including a qualification standard, employment test, or proxy) that discriminates against an applicant or employee or a class of applicants or employees on a basis protected under the FEHA, subject to any available defenses. Any evidence, or lack of evidence, of anti-bias testing or similar proactive efforts to avoid unlawful discrimination, including the quality, efficacy, recency, and scope of such efforts, the results of such testing or other efforts, and the response to the results, are relevant to any claims of discrimination or defenses thereto. The regulations incorporate restrictions on automated decision systems into the existing rules regarding discrimination of

all types, and also apply these restrictions to all applicable stages of employment and employment practices, with significant focus on pre-employment practices and employee selection practices. Importantly, the regulations require preservation and retention of automated decision system data dealing with any employment practice and affecting any employment benefit of any applicant or employee.[1]

2025 amendments to the CCPA regulations govern the use of automated decisionmaking technologies (ADMT) and impose requirements on employers' use of artificial intelligence and other automated tools in employment decisionmaking. Employers subject to the CCPA who use ADMT for employment-related decisions, without meaningful human involvement, must conduct detailed risk assessments, provide pre-use notices, and honor certain opt-out and access rights. Specifically, the regulations apply to ADMT used to make significant decisions about California residents, which include decisions that result in the provision or denial of employment or independent contracting, including hiring, assignment of work, compensation, promotion and demotion, and termination. The required risk assessment must evaluate whether the privacy risks of the ADMT outweigh its benefits to the consumer, the business, and other stakeholders. If the risks outweigh the benefits, the employer may not proceed with using the ADMT unless it can take steps to sufficiently mitigate the risks.

Employers can remove technologies from the ADMT rules entirely by ensuring that employment decisions are made with meaningful human involvement. For the human involvement to be adequate, the rules require that the decisionmaker (1) knows how to interpret the ADMT's outputs, (2) analyzes the ADMT's output and any other information relevant to the decision, and (3) has the authority to make or change the decision based on that analysis.[2]

Colorado. Enacted in 2025 and significantly amended in 2026, the Colorado Automated Decision-Making Technology Act regulates the use of automated decision-making technology (ADMT) to materially influence a consequential decision, which includes access to, eligibility for, selection for, or compensation for employment and imposes three primary compliance requirements on employers: pre-use notice, an adverse action process, and record retention.[3]

The Act splits liability for any potential algorithmic discrimination claim between developers (the entity that developed and supplied the ADMT tool or system) and deployers (the end user, such as an employer). Under the Act, both the developer and the deployer may be held liable for unlawful discrimination arising from use of covered ADMT, but only to the extent of their relative fault. The Act does not create joint and several liability, except to the extent permitted under existing law. The Act also invalidates contractual indemnification provisions against liability for the developer's or deployer's own acts or omissions related to the use of covered ADMT in violation of Colorado's antidiscrimination laws. However, this restriction does not apply if the covered ADMT was used in an unintended manner and the developer complied with its documentation obligations. The Act takes effect January 1, 2027.[4]

Connecticut. In 2026, the state enacted a law establishing requirements for and restrictions on an employer's use of AI in employment decisions. The requirements center on use of an automated employment decision tool that generates output that is a substantial factor used to make or materially influence an employment-related decision.

Like the Colorado AI Act, this law divides responsibilities between developers and deployers. Employers will typically be deployers, although a developer and a deployer may agree by contract to allocate certain compliance responsibilities between them. The law also creates notice and disclosure duties for employers that deploy automated employment decision tools in the state on or after October 1, 2027.[5]

Illinois. Like California, Illinois has multiple laws in effect that regulate how an employer may use AI tools. 2024 amendments to the Illinois Human Rights Act (IHRA) sets limitations on how an employer may use AI systems in making employment-related decisions. The IHRA is the state's general antidiscrimination statute and prohibits discrimination in employment on the basis of an applicant's or employee's membership in a protected class. The IHRA prohibits an employer from using AI functionality that has the effect of subjecting employees and applicants to discrimination on the basis of a protected class in connection with the following facets of the employment lifecycle: recruitment, hiring, promotion, renewal of employment, selection for training or apprenticeship, discharge, discipline, tenure, or the terms, privileges, or conditions of employment. An employer is also required to provide notice to applicants and employees that it is using AI for actions related to any of these purposes. Moreover, the Act specifically prohibits the use of zip codes as a proxy for protected classes when making employment-related decisions. While the law is ambiguous on the issue, the language implies that this refers to the use of zip codes as part of the information provided to an AI tool or process.[6]

The Illinois Artificial Intelligence Video Interview Act, enacted in 2019, applies to employers that ask applicants to record video interviews and use artificial intelligence analysis of the applicant-submitted videos. Under the Act, an employer must notify applicants that AI analysis tools will be used and obtain applicants' consent. If an employer does not obtain an applicant's consent to use artificial intelligence, the employer may not use artificial intelligence analysis to evaluate that applicant.

An employer that relies exclusively on a video interview's artificial intelligence analysis to determine whether an applicant will be selected for an in-person interview must collect and report the race and ethnicity of applicants who were not selected for an in-person interview, and the race and ethnicity of applicants who were hired to the Illinois Department of Commerce and Economic Opportunity annually by December 31st.[7]

New Jersey. Similar to California's FEHA regulations, the regulations implementing New Jersey's antidiscrimination statute, the Law Against Discrimination (NJLAD), were amended in 2025 in part to address the use of AI tools in employment. The AI provisions were part of broader amendments which codified the disparate impact theory of liability under the NJLAD. The regulations set forth several ways in which use of an artificial intelligence or automated decisionmaking tool (ADMT) may result in disparate impact liability. Significantly, the regulations suggest that an employer must test any ADMT it intends to use prior to its use, or ensure that the developer of the ADMT has adequately assessed the ADMT for potential bias. An employer's use of an ADMT that has not been adequately tested and shown to not adversely affect people in a protected class before its use may have a disparate impact on members of that protected class.[8]

New York City. Enacted in 2021, Local Law 144 regulates the use of automated employment decision tools in hiring and promotion decisions. The law requires that any such tool undergo an annual, independent bias audit conducted by an independent auditor to assess whether the tool's selection criteria result in disparate impact based on race, ethnicity or sex, and the employer must make a summary of the audit results publicly available. An employer using a covered AI tool must provide each job candidate (both internal or external) with ten business days' advance notice if an AI tool will be used in the decisionmaking process, listing the job qualifications and characteristics used to make the assessment. Further, the notice must specify the sources and types of data the AI tool uses, and the employer must provide a copy of its data retention policy. Candidates must be able to opt out and request an alternative selection process or accommodation.[9]

B. Restrictions on "Bossware"

California. The state's 2026 AI legislation package included a law that restricts AI-enabled workplace monitoring. The law, which takes effect January 1, 2027, prohibits an employer from using a workplace surveillance tool that uses AI to assess employees by (1) recognizing or making inferences or predictions about a person's emotional state; or (2) collecting neural data, defined as information that is generated by measuring the activity of an employee's central or peripheral nervous system, and that is not inferred from non-neural information.[10]

Connecticut and Maine. While these states' electronic monitoring restrictions do not expressly mention AI, the covered technology as defined under these laws would pull in AI-enabled tools. Connecticut's law requires an employer that engages in any type of electronic monitoring to give prior written notice to all employees who may be affected via notice at the time of hire as well as a workplace posting, informing them of the types of monitoring that may occur. Similarly, Maine's law provides that an employer cannot engage in electronic surveillance of employees unless the employer notifies the employees before beginning the use of surveillance.[11]

C. AI-Related Layoff Notice Requirements

California. As part of the package of AI-related bills enacted in 2026, the state enacted a new law amending the California Worker Adjustment and Retraining Notification law (Cal-WARN) to require layoff notices to identify whether the separation is due to displacement by an AI system or other automated technology. The amendments apply to a Cal-WARN-covered employer that is conducting a mass layoff, relocation, or termination caused in whole or in substantial part by an AI system or other automated technology replacing or automating employment positions. In addition to the information already required under Cal-WARN for the 60-day notice of mass layoff, relocation, or termination, the notice

for a technological displacement must also include information on the number and types of workers displaced by AI and the job functions that will be automated or otherwise performed by AI tools.[12]

Connecticut. Connecticut's law, also enacted in 2026, requires disclosure related to the impact of AI use on business restructuring. An employer that conducts layoffs that trigger the requirement to provide a layoff notice to the Connecticut Department of Labor under the federal Worker Adjustment and Retraining Notification (WARN) Act must also disclose to the Department whether the layoffs are related to the employer's use of AI or another technological change.[13]

New York. While this requirement is not currently captured in any statute or regulation, New York State Worker Adjustment and Retraining Notification (WARN) Act notices must disclose whether layoffs are related to the employer's use of AI. In her January 2025 State of the State plan, Governor Kathy Hochul announced that she would direct the New York Department of Labor (NYDOL) to amend the NY WARN regulations to require that notices contain this technology-related disclosure in order to provide the NYDOL with real data on the potential impact of technology on workers.[14] As of November 2026, the NY WARN regulations have not yet been amended to this effect, but shortly after Governor Hochul's announcement, the NYDOL's WARN Portal,[15] where employers are required to file their notices, was updated to require employers to make the disclosure on AI-related separations.

D. "Can't Blame AI": Laws Allocating Liability for AI Harm

California. The California Civil Code governs the prohibited and permitted civil defenses related to harm caused by artificial intelligence systems. The law provides that in an action against a defendant who developed, modified, or used AI that is alleged to have caused harm to the plaintiff, the defendant cannot assert that the AI autonomously caused the harm to the plaintiff. However, a defendant may still provide any other affirmative defense, including evidence relevant to causation or foreseeability. In addition, the defendant may present other evidence relevant to the comparative fault of any other person or entity.[16]

Connecticut. As part of its 2026 omnibus AI legislation, the state enacted a law amending the Connecticut Fair Employment Practices Act to provide that an employer's use of an automated employment decision tool is not a defense to a discrimination complaint. However, the Connecticut Commission on Human Rights and Opportunities or a reviewing court may consider evidence that the employer conducted anti-bias testing or undertook similar proactive efforts to prevent discrimination, including but not limited to the quality, efficacy, recency and scope of the testing or efforts, the results of the testing or efforts, and the response thereto. In short, bias testing will not immunize an employer from liability, but it may be weighed in the employer's favor.[17]

II. Overview of Global Regulatory Frameworks

The AI laws in place in the European Union, Canada, the United Kingdom, China, South Korea, and Brazil follow different models. The EU and South Korea have enacted comprehensive, risk-based AI statutes that differ considerably in scope and intensity. China regulates through layered, technology- and service-specific administrative rules alongside its data protection law. The United Kingdom, Canada, and Brazil have no AI-specific statute in force and rely primarily on data protection law, sectoral or provincial rules, and regulatory guidance.

The European Union has adopted the most developed horizontal regime. Under the EU Artificial Intelligence Act ("EU AI Act"), AI systems intended for recruitment or candidate selection, or for decisions concerning work-related relationships, promotion, termination, task allocation based on individual behavior or personal traits, or worker monitoring and evaluation, are generally classified as high-risk, subject to a narrow exemption that does not apply where the system profiles natural persons.[18] In practice, this exemption is therefore less widely applicable than some employers may expect. Where such a system automatically processes personal data for the purpose of evaluating certain personal aspects relating to individuals, such as their performance at work, reliability or behavior, that processing constitutes profiling, making the exemption unavailable. Following Regulation (EU) 2026/1744 ("Digital Omnibus on AI"), the requirements and operator obligations for these systems, including human oversight under Article 14, apply from December 2, 2027, rather than August 2, 2026.[19] Since February 2, 2025, however, AI systems inferring emotions in the workplace have been

prohibited, save for medical or safety reasons.[20] The General Data Protection Regulation (GDPR) separately regulates decisions based solely on automated processing that produce legal or similarly significant effects.[21] The Platform Work Directive adds algorithmic-management rules for digital labor platforms, to be transposed by December 2, 2026.[22]

South Korea's *Framework Act on the Development of Artificial Intelligence and the Establishment of a Foundation for Trustworthiness*, Act No. 20676 ("AI Basic Act"), has applied since January 22, 2026.[23] It regulates generative, high-impact, and high-performance AI and lists evaluations significantly affecting individuals' rights, such as hiring, among potentially high-impact uses.[24] The South Korean Ministry of Science and ICT (MSIT) has announced that it will refrain from imposing administrative fines (up to KRW 30 million, or approximately US $22,600) for at least one year; the underlying obligations nonetheless apply.

The remaining jurisdictions rely more heavily on privacy, public sector, employment, or activity-specific rules. The United Kingdom relies on a principles-based, regulator-led approach and the amended UK GDPR.[25] China combines its data protection law with rules on algorithmic recommendation, deep synthesis, and public-facing generative AI services, and issued policy guidance on AI agents in 2026.[26] Brazil's General Data Protection Law ("LGPD") is Brazil's principal enacted framework.

Appendix A Endnotes

1 *See* Cal. Code Regs. tit. 2, §§ 11008.1 – 11079; Grace L. Waddell, Joy C. Rosenquist, Alice Wang, and Niloy Ray, *California Approves Landmark AI Employment Regulations*, Littler ASAP (July 30, 2025).

2 *See* Cal. Code Regs. tit. 11, §§ 7000 – 7302; Zoe Argento, Philip Gordon, and Niloy Ray, *California's Long-Awaited Final Regulations on Automated Decisionmaking Create New Compliance Challenges for Employers*, Littler ASAP (September 26, 2025).

3 Colo. Rev. Stat. §§ 6-1-1701 through 6-1-1709.

4 Colo. Rev. Stat. §§ 6-1-1701 through 6-1-1709; Zoe Argento, Philip Gordon, and Alice Wang, *Colorado Amends its Artificial Intelligence Law, Substantially Reducing Obligations on Employers,* Littler ASAP (May 15, 2026).

5 Connecticut S.B. 5 (2026) (not yet codified in the Connecticut General Statutes); Allan G. King and Niloy Ray, *Connecticut Passes Law Significantly Regulating Use of AI in Employment,* Littler ASAP (May 14, 2026).

6 775 Ill. Comp. Stat. §§ 5/2-101 - 5/2-102; Michael Wilder, Niloy Ray, Elizabeth Hanford, and Jeffrey Iles, *Illinois Extends Statute of Limitations for Filing Discrimination Claims Under Illinois Human Rights Act, Adds Protected Classes, and Clarifies AI Use in Employment Decisions,* Littler ASAP (August 13, 2024).

7 820 Ill. Comp. Stat. § 42/20; Kwabena A. Appenteng, Philip L. Gordon, and Garry G. Mathiason, *Implementing Illinois' AI Video Interview Act: Five Steps Employers Can Take to Address Hidden Questions and Integrate Policies with Existing Employment Laws,* Littler ASAP (September 17, 2019).

8 N.J. Admin. Code §§ 13:16-1.3 - 13:16-4.1; Bradford J. Kelley and Tanner McCarron, *AI in the Garden State: New Guidance on Algorithmic Discrimination and the New Jersey Law Against Discrimination,* Littler ASAP (January 27, 2025).

9 Eli Freedberg, Niloy Ray, and Jim Paretti, *New York City Enacts Law that Hinders Use of Automated Tools in Hiring and Promotion Decisions*, Littler ASAP (December 28, 2021).

10 Cal. Lab. Code §§ 1580 – 1583; Joy Rosenquist, Zoe Argento, and Alice Wang, *California's 2026 AI Legislation Wave: New Laws Target AI Throughout the Employment Lifecycle*, Littler ASAP (October 14, 2026).

11 Me. Rev. Stat. tit. 26, § 620-A.

12 Cal. Lab. Code §§ 1400.5 – 1401; Shawn Clark, Robert Wilger, and Kerry Notestine, *California Amends its WARN Act (Again)*, Littler ASAP (Oct. 2, 2026).

13 Connecticut S.B. 5 (2026) (not yet codified in the Connecticut General Statutes); Allan G. King and Niloy Ray, *Connecticut Passes Law Significantly Regulating Use of AI in Employment,* Littler ASAP (May 14, 2026).

14 Gov. Kathy Hochul, *2025 State of the State Book* (January 14, 2025); Shawn Clark and Adriana Foreman, *Governor Hochul Seeks to Expand New York WARN Notice Requirements to Include AI Disclosures*, Littler ASAP (January 23, 2025).

15 https://nysdol.my.site.com/warnportal

16 Cal. Civ. Code § 1714.46.

17 Conn. Gen. Stat. § 46a-60; Allan G. King and Niloy Ray, *Connecticut Passes Law Significantly Regulating Use of AI in Employment*, Littler ASAP (May 14, 2026).

18 Article 6(2) in conjunction with Annex III, point 4, of Regulation (EU) 2024/1689 ("EU AI Act").

19 Article 113 of the EU AI Act, as amended by Regulation (EU) 2026/1744 of July 8, 2026, published in the Official Journal on July 24, 2026. For high-risk systems covered by Annex I (product legislation), the relevant date is August 2, 2028.

20 Article 5(1)(f) in conjunction with Article 113 of the EU AI Act.

21 Article 22 of Regulation (EU) 2016/679 ("GDPR").

22 Chapter III (Articles 7–15) and Article 29(1) of Directive (EU) 2024/2831 on improving working conditions in platform work ("Platform Work Directive").

23 Act No. 20676 of January 21, 2025, Addenda. An unofficial English overview is available at https://aibasicact.kr/.

24 Article 2(4) of the AI Basic Act; see also the high-impact AI determination guidelines issued by the Ministry of Science and ICT ("MSIT"), which list recruitment among the covered fields.

25 DSIT, *A pro-innovation approach to AI regulation: government response* (Feb. 6, 2024).

26 Article 20 of the Cybersecurity Law of the People's Republic of China, as amended on October 28, 2025; *see also* Interim Measures for the Management of Generative Artificial Intelligence Services (eff. Aug. 15, 2023).

Littler®

