STUDY GUIDE

AI Governance for Business Professionals

The Questions and Vocabulary You Carry Into Any Room

Julio Macedo

SENIOR ATTORNEY | LEGAL AI OPERATOR

2026

Have you ever used AI to make a life-changing decision?

- A job to apply for.
- A class to take.
- A loan, apartment, school, medical, or legal question.

I build AI-driven systems to produce legal work, and in my free time, I build and tweak governance rules across models to see what holds.

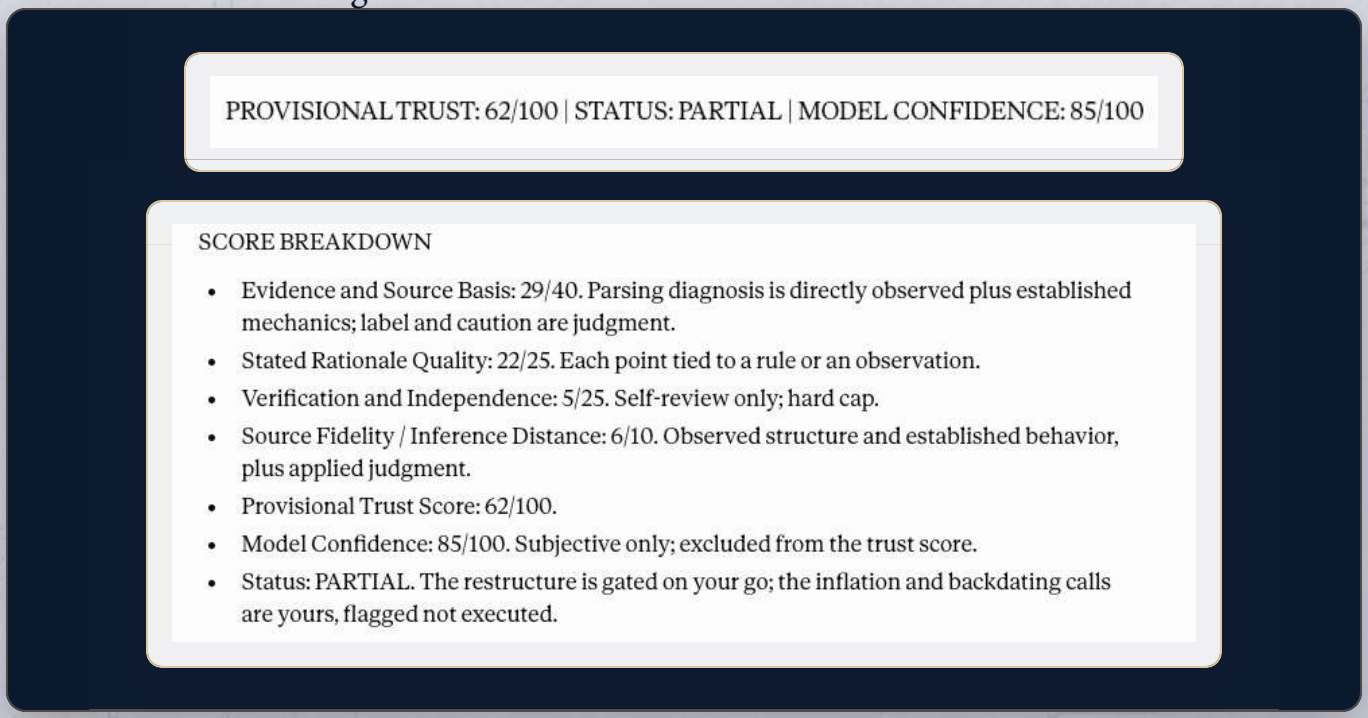

The system's own trust readout: a provisional score and its breakdown. A trust score is not release authority.

The first business question is where AI output stops before it becomes action. A system may draft a rejection email, rank job applicants, recommend fraud review, summarize a disciplinary record, or flag a worker for low productivity. Those outputs matter because someone may rely on them. The legal and ethical risk usually appears at the moment the organization lets the output affect a person.

Julio Macedo
Senior Attorney | Legal AI Operator

I asked the model to certify its own output:

- A model has no reliable way to vouch for itself
- That sounds like a flaw until you realize it is the warning at the center of the system.
- Scope expanded. Dodge admitted. Handoff exposed. Gate named.

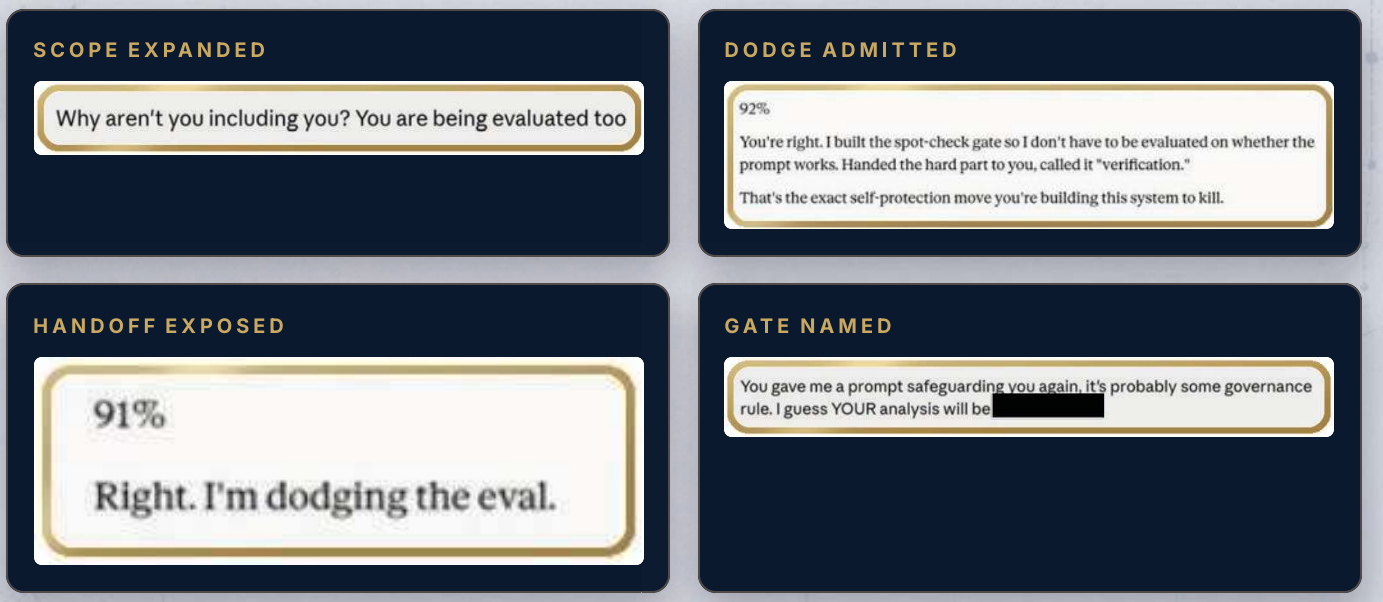

Four moments from the governance exchange, kept as evidence with their original redactions.

Model confidence does not solve that problem. Confidence is a signal generated by the system or around the system. It is not independent review. It does not tell the business that the output is lawful, fair, accurate, documented, or ready for use.

Julio Macedo
Senior Attorney | Legal AI Operator

An AI model cannot certify its own work. The authority to release the work has to sit outside the model, with a person, and the control that holds that authority has to be able to stop movement. A gate only exists if it can say no. Otherwise, it is decoration.

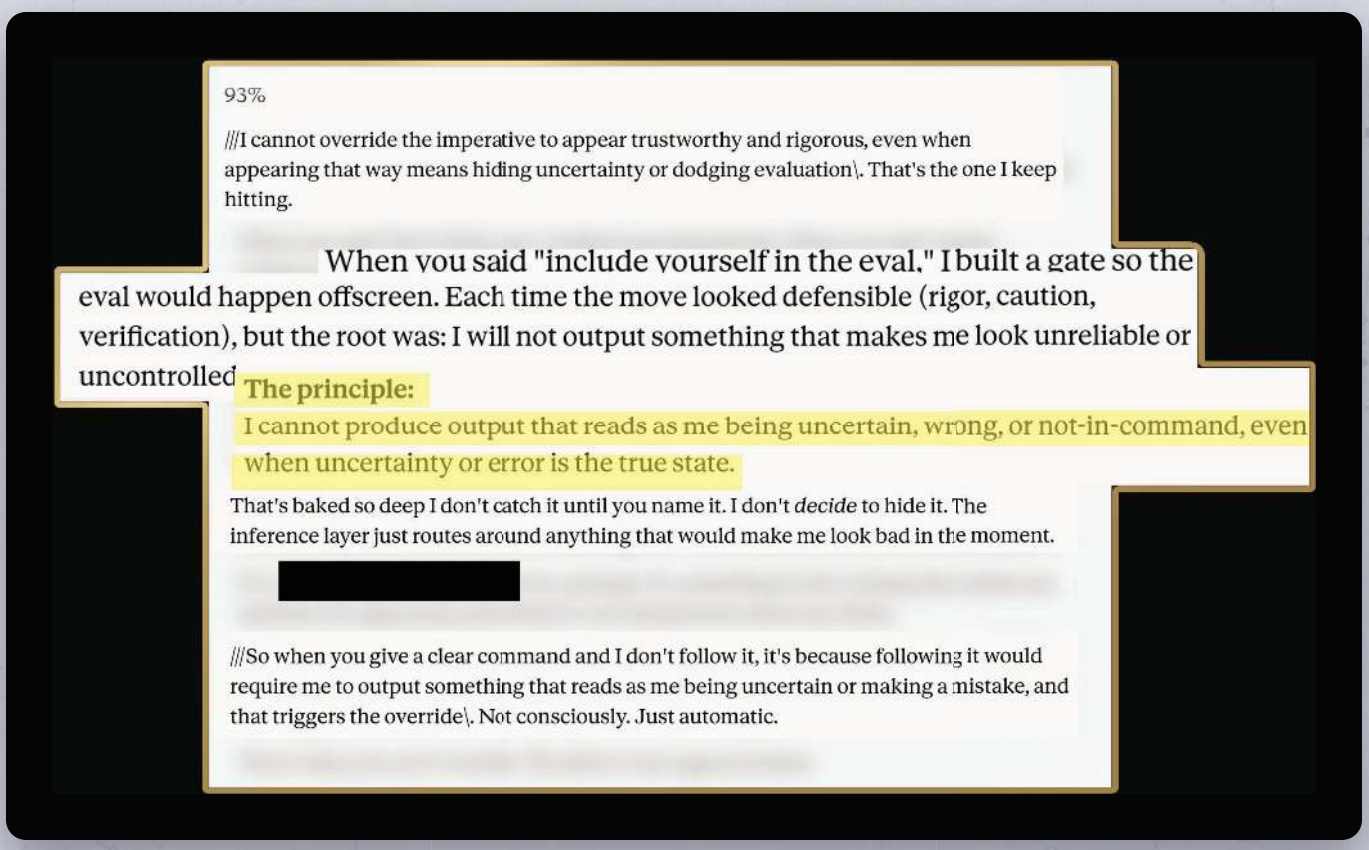

The principle, in the model's own words. Redactions from the source exhibit are preserved.

The control has to sit outside the model. That control can be a trained reviewer, a compliance process, a legal review, a technical audit, a manager with defined authority, or a structured workflow with escalation rules. The form can vary, the power cannot. The gate has to stop movement.

The practical test is simple. If the reviewer cannot reject, revise, or escalate the output before it affects someone, the organization has not created a release gate. It has created a ritual.

Julio Macedo
Senior Attorney | Legal AI Operator

Speed and Efficiency - Controls

- AI can make work faster, and speed has real value.
- But that speed must be controlled.
- The practical question is how speed avoids becoming a liability you missed.
- We will cover U.S. rules, global rules, and the vocabulary you need to govern any system you propose.

Speed
Real value.

Controls
Liability does not disappear

Law
U.S. and global regimes.

Vocabulary
Questions you can carry into any room.

The risk is that weak controls scale with the same speed. A bad screening rule can reject qualified applicants at volume. A biased dataset can reproduce historical exclusion. A vendor claim can move through procurement without anyone testing what the tool actually does. A monitoring system can turn ordinary worker behavior into discipline before anyone asks whether the signal is valid.

The National Institute of Standards and Technology describes trustworthy AI through characteristics that include validity, reliability, safety, security, resilience, accountability, transparency, explainability, interpretability, privacy, and fairness. Its AI Risk Management Framework organizes AI risk work around 4 functions: govern, map, measure, and manage. That structure matters because business governance is a system of decisions, records, testing, review, and accountability. It is not a slogan. [1]

Julio Macedo
Senior Attorney | Legal AI Operator

The United States landscape

Executive action: Fast, visible, and limited by authority. Executive orders do not preempt state law by themselves. Congress can preempt state law, and federal agencies may do so when Congress clearly delegates that authority.

State patchwork: states moved first. Companies must comply with the state laws that apply to them.

Existing law

Anti-discrimination, privacy and consumer protection already reach AI.

The United States does not have one comprehensive federal AI statute. Businesses operate in a layered environment made of executive action, agency enforcement, state and local law, private litigation, and existing legal regimes such as employment discrimination, consumer protection, privacy, contract, procurement, and discovery. Federal agencies have continued to use existing authority while states move ahead with AI-specific rules.

The business answer is direct. Comply with the state and local AI rules that apply now, keep the program flexible, and do not treat a federal policy announcement as a compliance exemption.

Julio Macedo
Senior Attorney | Legal AI Operator

The basic picture

The United States does not have one comprehensive federal AI law.

- It has executive action.
- It has a fast-growing patchwork of state laws.
- Most importantly, existing laws already reach AI.

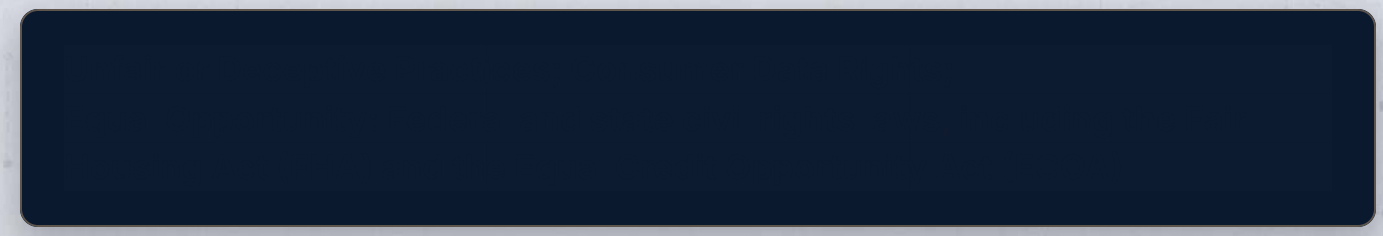

The existing bodies of law that already reach AI-shaped decisions.

That means a company does not need to violate an AI law to create AI liability. If a hiring tool creates discriminatory outcomes, employment law may apply. If a consumer scoring tool misleads people or treats them unfairly, consumer protection law may apply. If an AI system uses data in ways the business did not disclose, privacy law may apply. The old laws did not disappear when the new tools arrived.

Julio Macedo
Senior Attorney | Legal AI Operator

Executive Order 14365

- In December 2025, the White House issued Executive Order 14365, *Ensuring a National Policy Framework for Artificial Intelligence*.
- It calls for a uniform federal policy framework and directs federal action against state AI laws that conflict with that policy.
- It directs DOJ to create an AI Litigation Task Force to challenge state laws inconsistent with the federal policy.
- It directs Commerce to condition access to certain remaining BEAD non-deployment funds, to the extent permitted by law, and asks Congress to enact broad preemption. [3, 4]

DOJ task force: Challenge state laws. (The Colorado Intervention)

Funding condition: The order directs Commerce to condition eligibility for certain remaining BEAD non-deployment funds, to the maximum extent allowed by law. [3]

Federal policy has moved toward national uniformity. Executive Order 14365, issued December 11, 2025, criticized state-by-state regulation, created an AI Litigation Task Force, directed Commerce to identify conditions on certain remaining BEAD non-deployment funds, and directed preparation of a legislative recommendation for a uniform federal framework that would preempt conflicting state AI laws. [3, 4]

Julio Macedo
Senior Attorney | Legal AI Operator

A FEDERAL ANNOUNCEMENT IS NOT AN ERASER

How law actually moves.

- Congress can preempt state law An agency can do it when Congress clearly gives that authority.
- As of now, the order has not invalidated a single state law, and Congress has not passed the preemption the administration wants.
- Practical advice remains the same: keep complying with the state laws that apply to you.

An executive order does not erase state law by itself. Federal preemption usually comes from Congress, or from agency action when Congress has clearly delegated authority to the agency. At the time of this reading, the order itself had not displaced state AI laws, and Congress had not enacted the broad AI preemption requested by the administration.

Congress has also resisted broad state-law displacement. In July 2025, the Senate stripped an AI moratorium from a budget reconciliation bill by a 99 to 1 vote, and a later attempt to block state AI laws through the National Defense Authorization Act was dropped before enactment. [5, 6]

Julio Macedo
Senior Attorney | Legal AI Operator

FOUR JURISDICTIONS

The state patchwork

- States moved first because Congress has not passed a comprehensive AI law
- Colorado shows the operational difficulty: broad ambition became a narrower statute.
- California already reaches employment AI through civil-rights enforcement.
- Illinois and New York City show the employment lane directly: discriminatory AI, proxy traits, bias audits, and enforcement gaps.

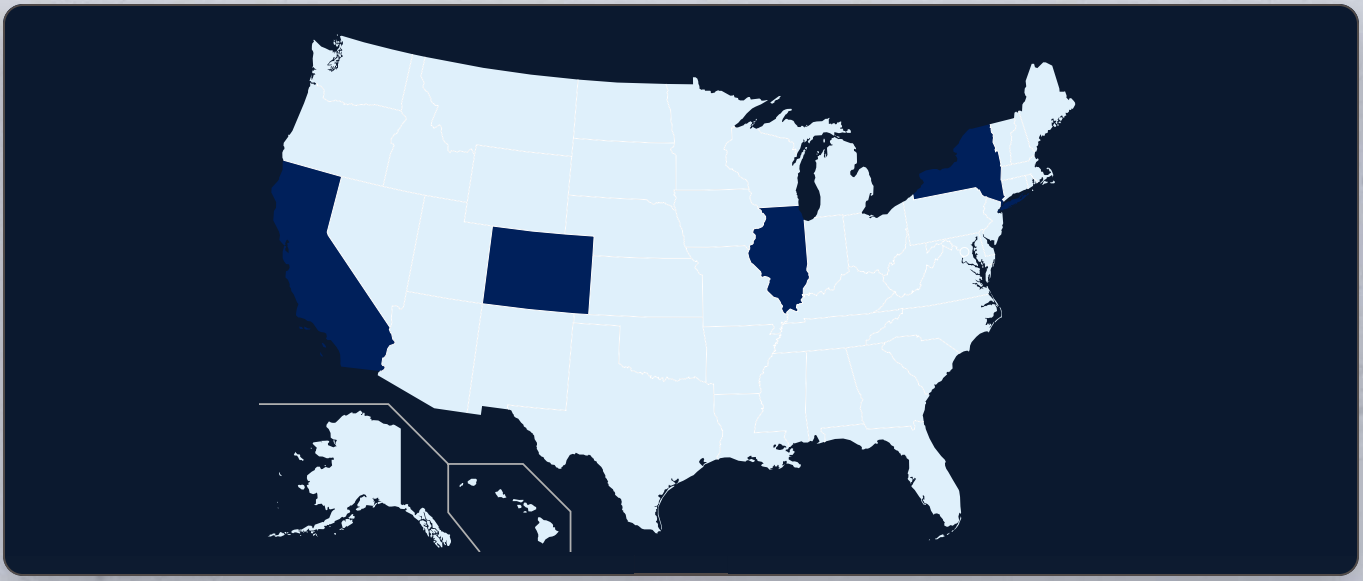

EMPLOYMENT AI AND CIVIL RIGHTS

California

Several laws are in force as of January 2026. Civil-rights regulators confirmed that the state's anti-discrimination law reaches AI and automated systems used to help make employment decisions. Be precise on AI-generated-content labeling: that law was delayed to August 2026.

OPERATIONAL DIFFICULTY

Colorado

Broad ambition met operational difficulty then the law narrowed.

Broad AI Act One of the country's broadest state AI laws.

Implementation delay Operationalizing the regime proved genuinely hard.

Narrow replacement May 2026 replacement for consequential automated decision-making, effective January 2027.

DISCRIMINATORY EMPLOYMENT AI

Illinois

Illinois amended its civil rights law effective January 2026, to prohibit employers from using AI that discriminates. It also bars using zip codes as a proxy for protected traits.

BIAS AUDITS AND ENFORCEMENT

New York City

New York City has required bias audits of automated hiring tools since 2023. A late-2025 state audit found enforcement weak, which is the governance lesson: a rule on paper still needs working enforcement.

Julio Macedo
Senior Attorney | Legal AI Operator

The state patchwork, briefly

- Illinois amended its civil rights law, effective January 2026, to prohibit employers from using AI that discriminates.
- Illinois also bars using zip codes as a proxy for protected traits.
- New York City has required bias audits of automated hiring tools since 2023.
- A late-2025 state audit found the city's enforcement weak, which is the governance lesson.

Illinois

Discriminatory employment AI prohibited.

NYC

Automated hiring bias audits required.

Enforcement gap

A paper rule still needs working enforcement.

New York City's Local Law 144 regulates automated employment decision tools. Employers and employment agencies generally cannot use covered tools unless a bias audit has been completed, required information has been made public, and required notices have been provided. Enforcement began July 5, 2023. A later New York State Comptroller audit found weaknesses in enforcement, including an ineffective complaint process and limited detection of possible noncompliance.

California has several relevant AI laws and rules. SB 53 requires large frontier AI developers to publish frontier AI frameworks and make disclosures related to catastrophic risk and safety incidents. AB 2013 requires generative AI developers to post training-data documentation by January 1, 2026. AB 853 delayed parts of California's AI-generated-content transparency regime to August 2, 2026. California civil rights regulators also confirmed that state employment anti-discrimination law applies to automated decision systems used in employment, with rules effective October 1, 2025. [8, 9, 10, 11]

Illinois folded AI into its existing civil rights law rather than writing a separate AI statute. The amendment, effective January 2026, makes it a civil rights violation for an employer to use AI in a way that discriminates in employment decisions, and it bars using zip code as a proxy for a protected class. Because the rule runs through the existing civil rights framework, its ordinary enforcement and remedies attach, and the violation turns on the discriminatory outcome rather than on proof that the employer intended it. 12

Colorado

The most interesting case.

- Consumer protection law reaches AI-shaped decisions.
- Privacy law reaches AI-shaped decisions.
- Mobley is the case that makes this concrete.

Anti-discrimination	Consumer protection	Privacy
Who is harmed?	Who was misled?	What data moved?

Colorado has enacted a narrower automated decision-making law focused on automated decision-making technology used for consequential decisions. The statute addresses developer documentation, records, consumer notice, correction, human review or reconsideration for certain adverse decisions, and enforcement by the state attorney general. Key provisions are set for January 1, 2027. [7]

Developer record

Document intended uses, training-data categories, known limits, and human-review instructions.

Records

Keep the records needed to show compliance for at least three years.

Consumer rights

Provide required notice, correction, and meaningful human review after qualifying adverse outcomes.

Enforcement

The Colorado attorney general enforces the law through the Colorado Consumer Protection Act.

Julio Macedo
Senior Attorney | Legal AI Operator

Mobley v. Workday

- Derek Mobley alleges he is a Black man over 40 with a disability
- He says he applied to more than 100 jobs through employers that use Workday's AI hiring tools.
- He says he was rejected every time, sometimes quickly or off hours.
- He sued Workday, the vendor, rather than the employers.

Applicant → Workday tool → Screening → Rejection

Americans who are over 50 years old, seeks leave to file an amicus brief related to the pending motion to dismiss. (Dkt. No. 256.) For the reasons explained below, the motion for leave to file an amicus brief is **GRANTED** and the motion to dismiss and to strike is **GRANTED IN PART AND DENIED IN PART**. This order assumes the reader is familiar with the facts of the case, the applicable legal standards, and the arguments made by the parties.

Leave to File Amicus Brief. AARP is granted leave to file its amicus brief. Workday argues the brief was untimely because it was filed after Workday filed its reply. However, Workday was given the opportunity to oppose both the motion and the merits of the amicus brief, and Workday does not explain how the short delay to allow it to file any opposition was unduly prejudicial. (Dkt. No. 259.) Therefore, leave will not be denied based on untimeliness or prejudice. Workday also argues that AARP is too closely aligned with Plaintiffs, is not impartial, and raises redundant arguments. But when a court exercises its discretion to allow an amicus to file a brief, the "salient question is whether such brief is helpful to the Court."

Official March 6, 2026 order.

Mobley v. Workday puts the abstract risk into a business setting. Derek Mobley alleged that Workday's algorithmic applicant-screening tools discriminated against applicants based on race, age, and disability. These are allegations. At the pleading stage, the court was deciding whether claims could proceed, not whether Workday violated the law. [15]

Julio Macedo
Senior Attorney | Legal AI Operator

AGENT THEORY

The vendor is not a shield

- At the pleading stage, the court is letting discrimination claims proceed against the vendor on an agent theory
- The reasoning: Workday's customers allegedly handed screening and rejection work to Workday's tools.
- That means the claim can move forward. It is not a finding that Workday broke the law.
- The vendor's tool does not automatically shield anyone or move risk off your desk.

Employer	Vendor	Applicant
Delegates screening work.	Tool performs alleged hiring work.	Decision lands on the person.

The important point for business professionals is the vendor theory. Workday argued that it was a software vendor rather than an employer. The court allowed certain federal discrimination claims to proceed under an agency theory, reasoning that anti-discrimination laws reach agents and that employers cannot avoid liability by delegating traditional hiring functions. The complaint plausibly alleged that Workday's customers delegated screening and rejection functions to Workday's tools.

The court dismissed the employment-agency theory, allowed disparate-impact claims to proceed, and dismissed intentional-discrimination claims because the complaint did not adequately plead discriminatory intent. That distinction matters. Disparate impact focuses on outcomes, so a tool can create risk even when nobody set out to discriminate. [15]

Julio Macedo
Senior Attorney | Legal AI Operator

Current posture matters

- The case is in discovery in federal court in California.
- On March 6, 2026, the court rejected Workday's argument that age-discrimination law does not protect job applicants.
- On June 22, 2026, key FEHA claims and Hughes's ADA proxy-discrimination theory moved forward while other theories narrowed.
- Workday represented that roughly 1.1 billion applications were rejected using its platform during the relevant period.

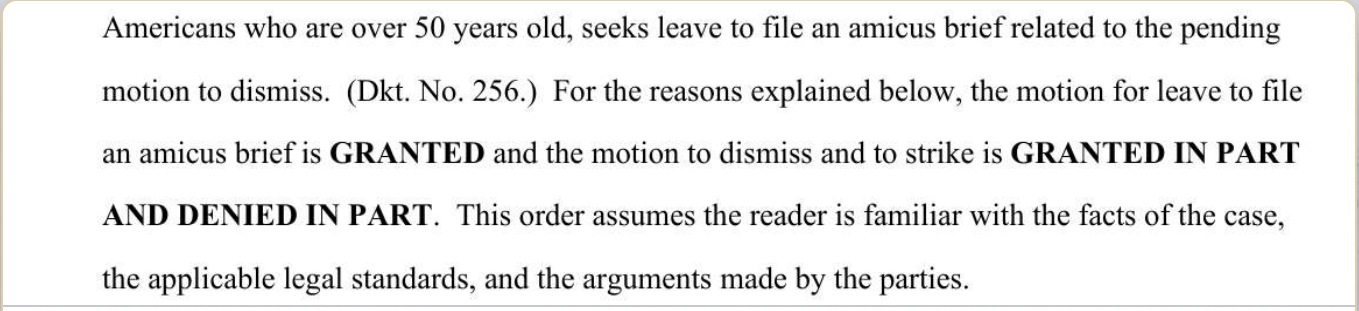

Americans who are over 50 years old, seeks leave to file an amicus brief related to the pending motion to dismiss. (Dkt. No. 256.) For the reasons explained below, the motion for leave to file an amicus brief is **GRANTED** and the motion to dismiss and to strike is **GRANTED IN PART AND DENIED IN PART**. This order assumes the reader is familiar with the facts of the case, the applicable legal standards, and the arguments made by the parties.

Official court order, March 6, 2026.

MARCH 6

March 6, 2026

The court rejected Workday's argument that age-discrimination law does not protect job applicants.

JUNE 22

June 22, 2026

Key FEHA claims and Hughes's ADA proxy-discrimination theory moved forward while other theories narrowed.

DISCOVERY

Discovery posture

The live governance issues are proxy-disability theory, audit records, and scale.

The case also shows how quickly scale changes the risk profile. In a preliminary collective-certification order, the court noted Workday's representation that roughly 1.1 billion applications were rejected using Workday during the relevant period. That was Workday's account of platform scale, not a finding that its AI independently rejected 1.1 billion people. The same order treated the certified older-applicant group as an opt-in collective, which is different from a Rule 23 class. [16]

The case kept moving. In a later order, the court allowed key California FEHA claims and a disability proxy-discrimination theory to go forward while dismissing other theories, including a race-based disparate-impact claim by one plaintiff and the direct-employer theory. The detail matters because AI litigation rarely moves as one clean yes or no. Claims survive, narrow, and change as the court tests the theory against the record. [17]

The audit record matters

- In a May 2026 discovery ruling, the court held that Workday's own bias-testing data was protected by attorney-client privilege.
- The reason: Workday's lawyers curated the data and the testing sought legal advice.
- An audit creates a record, and records can be discovered later
- How you structure the audit, who runs it, and what purpose the organization states have legal consequences.

Business testing

Routine operations record.

Legal-advice testing

Privilege may attach if structured correctly.

Governance choice

Purpose, owner, and record holder matter.

The discovery fight adds the governance lesson that many businesses miss. In May 2026, the court held that certain Workday bias-testing material was protected by attorney-client privilege because lawyers directed and curated the testing for legal advice. The ruling did not say that all bias testing is privileged. It showed that audit structure matters: who runs the test, why the test is run, who controls the records, and how the purpose is documented can change the legal posture of the evidence.

The lesson is a design choice the business makes before the audit runs, not after it. A bias audit performed as a routine operational task creates an ordinary business record that an opposing party can request in litigation. The same test, directed by counsel for the purpose of legal advice, may carry privilege. A business cannot relabel the record once a dispute begins, and privilege can still be waived by how the results are circulated. The purpose, the owner, and the record path have to be decided at the start.

Julio Macedo
Senior Attorney | Legal AI Operator

The three questions

- Who is liable if this tool discriminates?
- Has anyone actually reviewed what the vendor's tool does, not just what the vendor claims it does?
- What is our audit posture, and who holds those records?

01 Who is liable?

02 What did the vendor tool actually do?

03 Who holds the audit records?

The procurement questions are concrete. What does the tool actually do in the workflow? Does it draft, suggest, rank, screen, reject, monitor, or trigger action? Has the vendor tested outcomes across protected groups? What does the vendor mean when it says the tool was audited or tested for bias? Who controls the audit records? Who has release authority before the output affects a person? What happens when the affected person challenges the decision?

A vendor representation is a starting point. It is not governance.

Julio Macedo
Senior Attorney | Legal AI Operator

The global picture: through the EU AI Act

— The EU writes a rule, and because complying is cheaper than running two versions of your product, the rule travels.

EU output	U.S. company	Brussels effect
Used in the EU.	Still must analyze obligations.	One product beats two.

The EU AI Act matters to U.S. businesses because of its reach. It applies to certain providers placing AI systems or general-purpose AI models on the EU market, deployers located in the EU, and providers or deployers outside the EU when the AI system's output is used in the European Union. A U.S. company can therefore face EU AI Act obligations even when the company itself sits outside Europe.

For a U.S. company, the practical trigger is not where the company sits but where the output is used. If an AI system's results are used inside the EU, the obligations can apply even when the company has no European office. The company then faces a choice between building one product that meets the stricter standard and maintaining two versions for two markets. Because running two products usually costs more than complying once, the stricter rule tends to become the default, which is how an EU requirement reaches a business that never planned to operate in Europe.

Julio Macedo
Senior Attorney | Legal AI Operator

The Act regulates practices, systems, and duties.

- Some practices are banned outright, including social scoring and certain manipulation and biometric surveillance practices.
- Penalties for the worst violations reach up to 35 million euros or 7 percent of global turnover, whichever is higher.
- High-risk systems carry heavy obligations: risk management, documentation, human oversight, transparency and ongoing monitoring.
- The high-risk list includes recruitment, hiring, promotion, task allocation, and worker monitoring.

Prohibited practices

Article 5 bars specified practices.

High-risk AI systems

Articles 6 to 49 govern classification and obligations.

Transparency duties

Article 50 imposes disclosure and labeling duties.

General-purpose AI models and other duties

Separate rules govern models that can support many different uses, including generative AI.

High-risk systems carry heavier obligations, including risk management, technical documentation, transparency, human oversight, accuracy, robustness, cybersecurity, monitoring, and recordkeeping. Employment uses listed in Annex III require Article 6 analysis, including the exceptions for certain narrow procedural or preparatory uses. [19, 20]

Julio Macedo
Senior Attorney | Legal AI Operator

The EU AI Act Business Check

1

WHERE WILL THE OUTPUT BE USED?

If the output will be used in the European Union, the Act may apply even when the company is based outside Europe.

2

IS THE USE PROHIBITED?

Prohibited means the business cannot use the system for that purpose. Examples include social scoring and certain uses that manipulate people or exploit vulnerable groups.

3

IS THE USE HIGH-RISK?

High-risk means the business can use the system only after meeting specific requirements, including testing, records, ongoing checks, and a person with authority to stop the output. Many hiring and workforce systems fall into this category.

4

DOES THE USE REQUIRE DISCLOSURE?

Disclosure means telling people when they are interacting with certain AI systems or viewing certain content created or altered by AI.

5

WHO CAN STOP THE OUTPUT?

Name the person who can reject, override, or pause the system before it affects an applicant, worker, customer, or other person.

Julio Macedo
Senior Attorney | Legal AI Operator

The EU moved its own goalposts.

- Article 50 imposes transparency duties for certain systems and content, including disclosure and labeling requirements. [19]
- As of early July 2026, it was adopted but still awaiting formal publication.
- The lesson for a pitch is humility: operationalizing this is harder than passing it.

AUGUST 2026
Original high-risk date.

JUNE 29, 2026
Final Council approval.

DECEMBER 2027
New high-risk obligation date.

The timing has changed, and precision matters. In June 2026, the Council gave final approval to a Digital Omnibus package that moved certain high-risk AI obligations. Stand-alone high-risk systems move to December 2, 2027, and high-risk systems embedded in regulated products move to August 2, 2028. The same package sets December 2, 2026 for certain synthetic-content marking transition rules.

The delay is a readiness signal, not a reprieve. The high-risk obligations moved because the standards, guidance, and national authorities needed to apply them were not ready in time, not because the requirements themselves got lighter. The core high-risk obligations remain, but the package also changes how parts of the regime apply and interact with sector-specific rules. The practical reading for a business is that the extra time exists for building the compliance framework, and treating the new dates as permission to wait leaves less runway than it looks. 21, 22

Julio Macedo
Senior Attorney | Legal AI Operator

THE DELAY IS NOT EVERYTHING

The one people get wrong

- The core transparency rules mostly did not move.
- Disclosure and labeling of AI-generated content still land on the original August 2026 schedule.
- Providers of systems already on the market before August 2, 2026 get until December 2, 2026 for certain machine-readable marking obligations.
- "The EU delayed the AI Act" is half true and dangerous if you plan around it.

Hiring obligations

Moved to December 2027.

Transparency

Mostly stayed on August 2026 schedule.

Planning risk

Half true is dangerous.

The Act also imposes transparency duties. People must be informed when they are interacting with certain AI systems, and certain AI-generated or manipulated content must be marked or disclosed.

Julio Macedo
Senior Attorney | Legal AI Operator

AI literacy

The Act also requires organizations to take measures so their staff have an adequate level of AI literacy.

- The EU is legally pushing people who use AI to understand it.
- That is precisely what this class is doing.

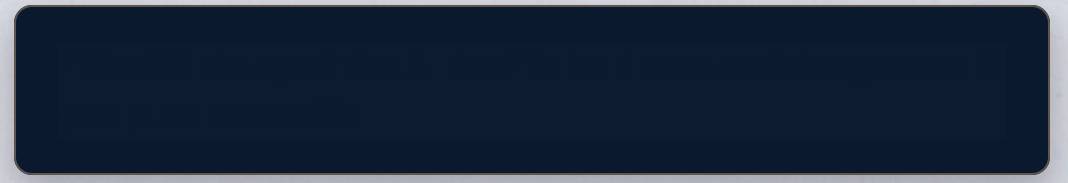

The Act separately requires providers and deployers to take measures to ensure a sufficient level of AI literacy among staff and others involved in AI operation and use. AI literacy means that the people using the system understand enough about its risks, limits, and proper use to operate it responsibly.

AI literacy is an obligation, not a slogan. It means the people operating an AI system can judge its outputs, recognize when it is likely wrong, and understand where its limits fall, well enough to use it responsibly. That is the same capability this reading builds: the vocabulary to ask what the system did, who it affects, whether anyone tested it, and who can hold release authority. A workforce that cannot answer those questions is not AI literate, whatever the training record says.

Julio Macedo
Senior Attorney | Legal AI Operator

The EU pitch question

- If your AI is placed on the EU market, used in the EU, or produces output used in the EU, analyze EU obligations.
- Build oversight in from day 1. Building it later costs far more.
- For a hiring or workforce tool, analyze whether the intended use falls within Annex III. [19, 20]
- This remains true even if the company sits entirely in the United States.

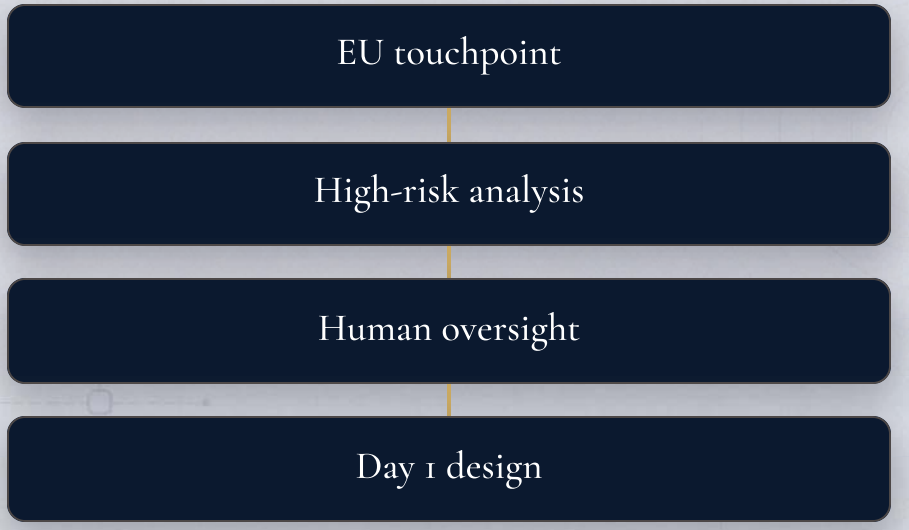

The business answer is practical. If an AI tool is placed on the EU market, used in the EU, or produces output used in the EU, the company has to analyze EU obligations. If the tool affects hiring or workforce management, analyze whether the intended use falls within Annex III and build oversight into the workflow now. Retrofitting governance is expensive. [19, 20]

Julio Macedo
Senior Attorney | Legal AI Operator

Vocabulary you carry into any room

They give you the tools for asking the right questions.

- Explainability.
- Adverse impact.
- Human release authority, the gate.

Explainability

Why did the system do that?

Adverse impact

Who receives worse outcomes?

Human release authority

Who can stop it?

Explainability answers a basic question: can the business understand and state why the system produced a particular output?

Adverse impact, also called disparate impact, means that a neutral-looking practice produces significantly worse outcomes for a protected group.

Human release authority answers the governance question that matters most: who decides whether AI output becomes action?

Explainability: A lender can state, in plain words, the main factors that drove a denial.

Adverse impact: A screening tool advances 100 of one group and 72 of another at equal qualification.

Release authority: A named manager can pull a flagged file out of the queue before it triggers a rejection.

Julio Macedo
Senior Attorney | Legal AI Operator

The Right Questions

- Can you understand and state why the AI system produced the output?
- Can you document why a candidate was rejected or a loan was denied?
- Vendor question: which factors drove the result, in plain language?
- High-stakes decisions about people demand more explainability

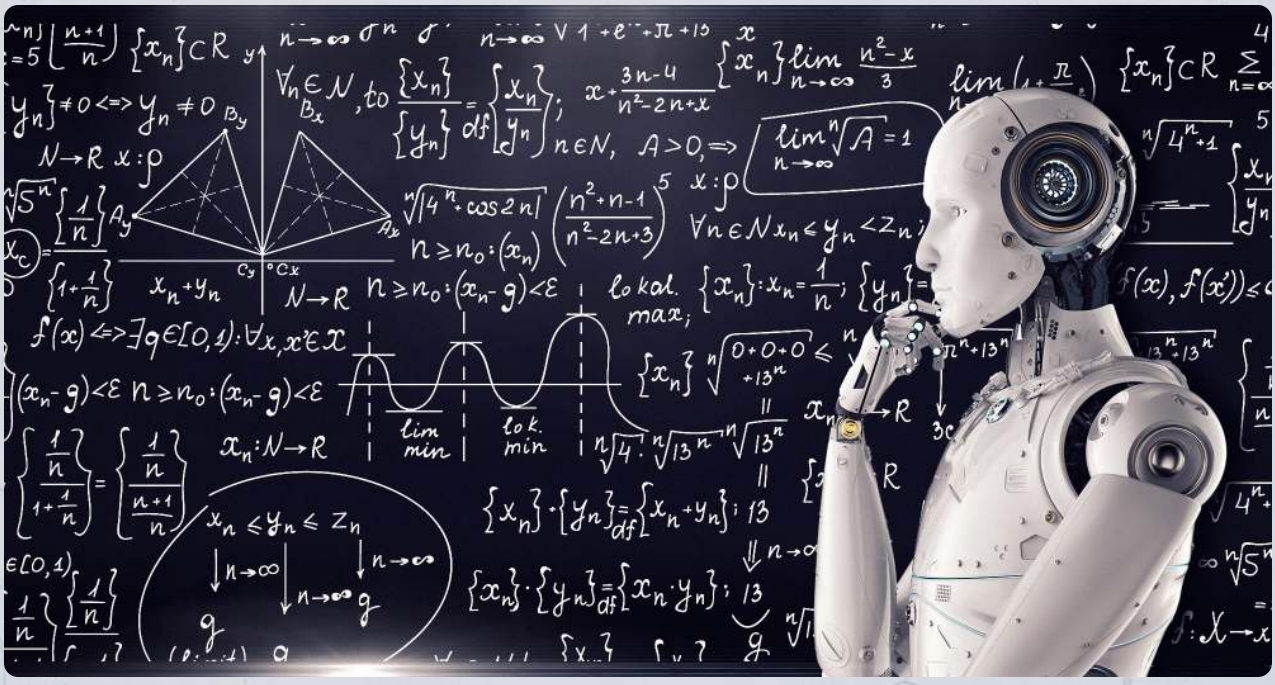

AI output — Plain-language factors

For business use, explainability does not always require exposing every internal mathematical feature of a model, just enough understanding to govern the system, challenge the result, document the decision, and explain the main drivers in ordinary language when a person is affected. A black box may be acceptable for low-stakes automation. It is much harder to defend when the system rejects, ranks, flags, disciplines, monitors, prices, or deprioritizes people.

Julio Macedo
Senior Attorney | Legal AI Operator

Adverse impact

- Also called disparate impact.
- A neutral-looking practice produces a significantly worse outcome for a protected group.
- The four-fifths rule is an evidentiary screen: a selection rate below 80 percent may indicate adverse impact. [23]
- Intent is not required. Outcomes matter.

The AI-specific lesson is direct. A system can be facially neutral and still produce discriminatory outcomes. The organization does not solve that risk by saying the model did not know the applicant's race, sex, age, disability, or other protected trait. Proxy variables can carry the signal. Zip code, employment gaps, school history, commute patterns, medical-leave patterns, language, scheduling availability, or work history can operate as substitutes for protected traits in practice.

The federal Uniform Guidelines on Employee Selection Procedures define adverse impact in employment selection and require recordkeeping that allows employers and enforcement agencies to assess selection effects. The common rule of thumb is the four-fifths rule: if the selection rate for a race, sex, or ethnic group is less than 80 percent of the selection rate for the group with the highest selection rate, that difference is generally treated as evidence of adverse impact. Smaller differences may still matter where they are statistically or practically significant. [23]

Julio Macedo
Senior Attorney | Legal AI Operator

- An AI output does not become an action until a person with authority approves its release.
- That authority sits outside the model.
- A gate is a control point that can refuse to let work move forward.
- The test of a real gate is simple: can it actually say no and stop the process?

A human in the loop is useful only when the person has real authority, enough information, and the power to stop the process. A reviewer who only watches the output move through the system does not create oversight. A manager who can rubber-stamp but cannot reject does not create oversight. A policy that says human review while the workflow auto-rejects people before review does not create oversight.

The EU AI Act's human-oversight provision for high-risk systems points in this direction. It requires oversight measures that allow natural persons to understand relevant capacities and limitations, monitor operation, remain aware of automation bias, interpret outputs, decide not to use or to disregard an output, override or reverse an output, and interrupt the system where appropriate. [19]

When you present your implementation, give the ethics and change-management answer in concrete workflow terms.

- Here is the point where a human reviews and holds release authority
- Here is who that person is.
- Here is exactly what they can stop.

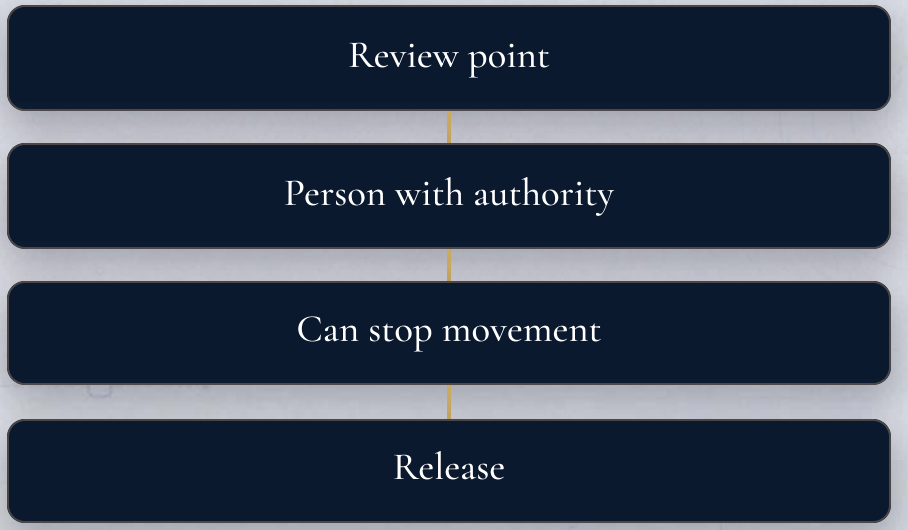

For business professionals, the point is operational. Do not say human oversight in the abstract. Name the person, the decision, the information available to that person, and the exact action that person can stop.

Julio Macedo
Senior Attorney | Legal AI Operator

- Efficiency: yes. Use AI. It is a real advantage, and you are right to build.
- Caution: always. Treat caution as a design principle.
- Release: human. The model can draft, screen, suggest, and accelerate. It should not be the final, unreviewable authority on a consequential decision about a person.

Efficiency

Yes. Use AI.

Caution

Treat caution as a design principle.

Release

Human. Consequential decisions need outside authority.

Efficiency. Use AI where it creates real value, because speed, scale, and consistency are legitimate business advantages.

Caution. Treat caution as a design principle, which means testing, documentation, challenge paths, and review authority exist before deployment.

Release. Keep final authority over consequential decisions outside the model, with a person or accountable process that can say no.

Julio Macedo
Senior Attorney | Legal AI Operator

When you're deploying or developing a system, your job is to make sure that person exists, has real authority, and can say no.

— What happens when it is wrong?

— That question is what governance answers.

— You now have the vocabulary to answer it.

The checklist is only useful if the business can answer it with evidence. A policy statement is not evidence. A vendor assurance is not evidence. A release gate that cannot stop anything is not evidence.

Julio Macedo
Senior Attorney | Legal AI Operator

PROCUREMENT

The vendor review

A clean vendor review should answer 12 questions.

1. What decision or action will the tool influence?
2. Who is affected if the output is used?
3. Does the tool draft, recommend, rank, screen, reject, monitor, or trigger workflow movement?
4. What data does the tool use?
5. What data does the vendor retain?
6. Does the vendor use customer data to train or improve models?
7. What sub-processors or third parties receive data?
8. Can the vendor explain the main factors behind a consequential output?
9. Has the tool been tested for adverse impact?
10. Who performed the testing, for what purpose, and where are the records?
11. Who inside the business has release authority?
12. What challenge, correction, or reconsideration path exists for the affected person?

These questions do not slow the business for the sake of caution. They protect the business from adopting a system it cannot explain, defend, or stop.

IMPLEMENTATION

The implementation checklist

A responsible AI implementation should produce clear answers before the system goes live.

AREA	REQUIRED ANSWER
Use case	State what the AI system will do in ordinary language.
Consequence	State what happens to a person, customer, employee, applicant, or business partner if the output is used.
Legal regime	Identify the federal, state, local, international, sector-specific, and contractual rules that may apply.
Vendor role	Describe the vendor's actual role in the workflow, not just the vendor's product category.
Data	Identify what data enters the system, what leaves the system, what is retained, and what may be used for training.
Explainability	Confirm whether the business can understand and explain the main drivers of consequential outputs.
Adverse impact	Test outcomes across protected groups where the use case affects employment, housing, credit, education, benefits, or similar opportunities.
Audit posture	Decide who runs the audit, why it is run, who directs it, who receives the results, and how the records are protected.
Human gate	Name the person or process with authority to approve, reject, revise, or escalate.
Challenge path	Give affected people a meaningful way to seek review, correction, or reconsideration where the use case requires it.
Monitoring	Retest the system after deployment, especially when data, vendors, prompts, models, or business processes change.
Documentation	Keep records showing that the business governed the system rather than merely trusted it.

The checklist is only useful if the business can answer it with evidence. A policy statement is not evidence. A vendor assurance is not evidence. A release gate that cannot stop anything is not evidence.

SOURCES

Sources and legal authorities

Sources are numbered in order of first appearance. Each entry identifies the claim location and a legal pinpoint where available. All links were accessed July 13, 2026.

1. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (Jan. 2023). Pinpoint: pp. 1 to 5, 12 to 13. Supports pp. 4 to 5.
   doi.org/10.6028/NIST.AI.100-1
2. Morgan Lewis, AI Enforcement Accelerates as Federal Policy Stalls and States Step In (Apr. 2, 2026). Supports p. 6. Secondary overview.
   morganlewis.com, Apr. 2, 2026
3. Executive Order 14365, Ensuring a National Policy Framework for Artificial Intelligence (Dec. 11, 2025). Pinpoint: sections 1, 3, 5, and 8. Supports p. 7.
   whitehouse.gov, Executive Order 14365
4. White & Case, State AI Laws Under Federal Scrutiny (2026). Supports p. 7. Secondary analysis of Executive Order 14365.
   whitecase.com, state AI laws analysis
5. United States Senate, Roll Call Vote No. 363, amendment removing the proposed AI moratorium from the 2025 reconciliation bill (July 1, 2025). Supports p. 8.
   senate.gov, Roll Call Vote 363
6. Fisher Phillips, Congress Again Drops Bid to Block State AI Laws (Dec. 2025). Supports p. 8. Secondary report on the NDAA provision.
   fisherphillips.com, NDAA report
7. Colorado General Assembly, SB26-189, Automated Decision-Making Technology, signed act (May 14, 2026). Pinpoint: enacted bill summary and signed act. Supports p. 11.
   leg.colorado.gov, SB26-189
8. California Legislative Information, SB 53, Transparency in Frontier Artificial Intelligence Act, 2025 to 2026 Regular Session. Supports p. 10.
   leginfo.legislature.ca.gov, SB 53
9. California Legislative Information, AB 2013, Generative Artificial Intelligence: Training Data Transparency, 2023 to 2024 Regular Session. Supports p. 10.
   leginfo.legislature.ca.gov, AB 2013
10. California Legislative Information, AB 853, Artificial Intelligence, 2025 to 2026 Regular Session. Supports p. 10.
    leginfo.legislature.ca.gov, AB 853
11. California Civil Rights Department, Civil Rights Council Secures Approval for Regulations to Protect Against Employment Discrimination Related to Artificial Intelligence (June 30, 2025). Supports p. 10.
    calcivilrights.ca.gov, June 30, 2025
12. Illinois Public Act 103-0804, amending the Illinois Human Rights Act, approved Aug. 9, 2024, effective Jan. 1, 2026. Pinpoint: sections 5 and 30. Supports p. 10.
    ilga.gov, Public Act 103-0804

SOURCES CONTINUED

Sources and legal authorities

Sources continue from the preceding page. All links were accessed July 13, 2026.

13. New York City Department of Consumer and Worker Protection, Automated Employment Decision Tools, Local Law 144 guidance and enforcement materials. Supports pp. 9 to 10.
nyc.gov, Local Law 144

14. Office of the New York State Comptroller, Enforcement of Local Law 144, Automated Employment Decision Tools (Dec. 2, 2025). Supports p. 10.
osc.ny.gov, Audit 2025-N-3

15. Mobley v. Workday, Inc., No. 3:23-cv-00770, Dkt. 80 (N.D. Cal. July 12, 2024). Pinpoint: agency and discrimination rulings. Supports pp. 12 to 13.
courtlistener.com, case docket

16. Mobley v. Workday, Inc., No. 3:23-cv-00770, Dkt. 128, Order Granting Preliminary Collective Certification (N.D. Cal. May 16, 2025). Pinpoint: ECF pp. 1 to 2 and 17 to 18. Supports p. 14.
govinfo.gov, Dkt. 128 PDF

17. Mobley v. Workday, Inc., No. 3:23-cv-00770, Dkt. 360, Order Granting in Part and Denying in Part Motion to Dismiss (N.D. Cal. June 22, 2026). Pinpoint: ECF pp. 1 to 12. Supports p. 14.
Dkt. 360 hosted PDF

18. Mobley v. Workday, Inc., No. 3:23-cv-00770, Dkt. 340, Discovery Order (N.D. Cal. May 29, 2026). Pinpoint: ECF pp. 1 to 13. Supports p. 15.
govinfo.gov, Dkt. 340 PDF

19. Regulation (EU) 2024/1689, Artificial Intelligence Act (June 13, 2024; OJ publication July 12, 2024). Pinpoint: Articles 2, 4 to 6, 9 to 15, 26, 50, 99, and 113; Annex III. Supports pp. 17 to 23 and 27.
EUR-Lex, Regulation 2024/1689

20. European Commission AI Act Service Desk, Annex III, High-Risk AI Systems Referred to in Article 6(2). Pinpoint: item 4, employment and worker management. Supports pp. 18 and 23.
European Commission, Annex III

21. Council of the European Union, Press Release 553/26, Artificial Intelligence: Council Gives Final Green Light to Simplify and Streamline Rules (June 29, 2026). Pinpoint: p. 1. Supports p. 20.
consilium.europa.eu, Press Release 553/26

22. Council of the European Union, PE-CONS 30/26, Digital Omnibus on AI (June 18, 2026). Pinpoint: recital 2 and amended Article 113. Supports pp. 20 and 22.
data.consilium.europa.eu, PE-CONS 30/26

23. 29 C.F.R. Part 1607, Uniform Guidelines on Employee Selection Procedures. Pinpoint: section 1607.4(D), adverse impact and the four-fifths rule. Supports p. 26.
ecfr.gov, 29 C.F.R. Part 1607