

BROUGHT TO YOU BY THE TEAM GUIDING HR THROUGH THE AI SHIFT

AI in HR WEEK

Powered by **AAIM**



DAY 3 · WEDNESDAY, JULY 22

Compliance and Governance

Your plain-language guide to what's changed — and how to keep AI between the guardrails as it gets closer to real decisions.

Compliance can feel like the heavy part of the week — the place where the fun of AI runs into a wall of rules. We see it differently.

Knowing the rules of the road is exactly what lets you drive with confidence.

The landscape has grown up fast this year, so today we'll walk through what's actually changed since last time, and a simple way to tell when an everyday AI task is asking for a little more care. No fear, no legalese — just an honest, HR-to-HR translation you can use right away. Let's get into it!



Rules of the Road

Picture a highway for a second.

Off to your right is a metal guardrail. You barely notice it; it isn't slowing you down. But it's part of why you can take that curve at full speed with confidence. The road signs work the same way — speed limit, curve ahead, merge coming up. You never memorized the traffic code to read them. You just know the rules of the road well enough to drive with confidence.

The same car, on the same road, can be driven with confidence or with a lot of anxiety. The difference is a driver who knows the rules and reads the signs. AI compliance is much the same. The guardrails and signs aren't there to stop you from using AI; they're what let you use it with confidence and speed, without white knuckles. The goal today is to help you read the few signs that are now posted.



So what's the first sign to read? It's a shift in how you think about the rules.

It's tempting to assume compliance is about which tool you use — that some apps are “safe” and others aren't. That's not where the rules live. What they actually care about is what AI helped you decide, and whether a real person was affected by it. Using AI to draft an internal reminder? No one's worried about that. Using it to help decide who gets hired, promoted, or disciplined? That's a different matter entirely.

One idea sits underneath everything today.

The closer AI gets to a decision about a person, the more care it deserves. The question to carry isn't “are we using AI?” — it's “what did AI help decide, and what happened to a person because of it?”



A NOTE FROM AI GURU LAUREN PATEL

If you were with us in 2024 and 2025, you'll notice this topic has grown a lot — that's not us making it complicated, it's the landscape maturing in real time, which is a good sign. My promise today: plain terms you can use, the deeper resources when you want them, and an honest “still being figured out” where that's true. Let's read the signs together.

What's Changed Since 2025

Last year we talked about the AI employment laws that were coming. This year, a few are already in effect, more have been signed with 2027 deadlines, and several states are still writing theirs. Here are a few highlights of where we stand today:

● LIVE NOW

Illinois

Since Jan 1, 2026

No AI in employment decisions in a way that has a discriminatory effect; employees must be told when it's used. Bars ZIP codes as proxies, and defines "AI" broadly — even everyday generative tools. (HB 3773)

California

Since Oct 1, 2025

Anti-discrimination law applies to automated systems; records kept 4+ years; "fun" AI assessments can cross into unlawful disability inquiries.

New York City

Since 2023

Automated hiring tools need a bias audit, a public summary, and candidate notice — where your applicant sits can matter as much as where your office sits. (Local Law 144)

○ COMING SOON

California — ADMT

Broader automated-decision rules finalized; deadline for AI in "significant decisions" is Jan 1, 2027.

Colorado

Covers tools that materially influence consequential decisions; effective Jan 1, 2027.

Connecticut

New disclosure and notice duties for covered employment AI; effective Oct 1, 2027.

Federal

Still no single federal AI employment law. A national framework was floated in early 2026 but needs Congress. The federal anti-discrimination laws you already follow still apply.

Your state isn't the whole story

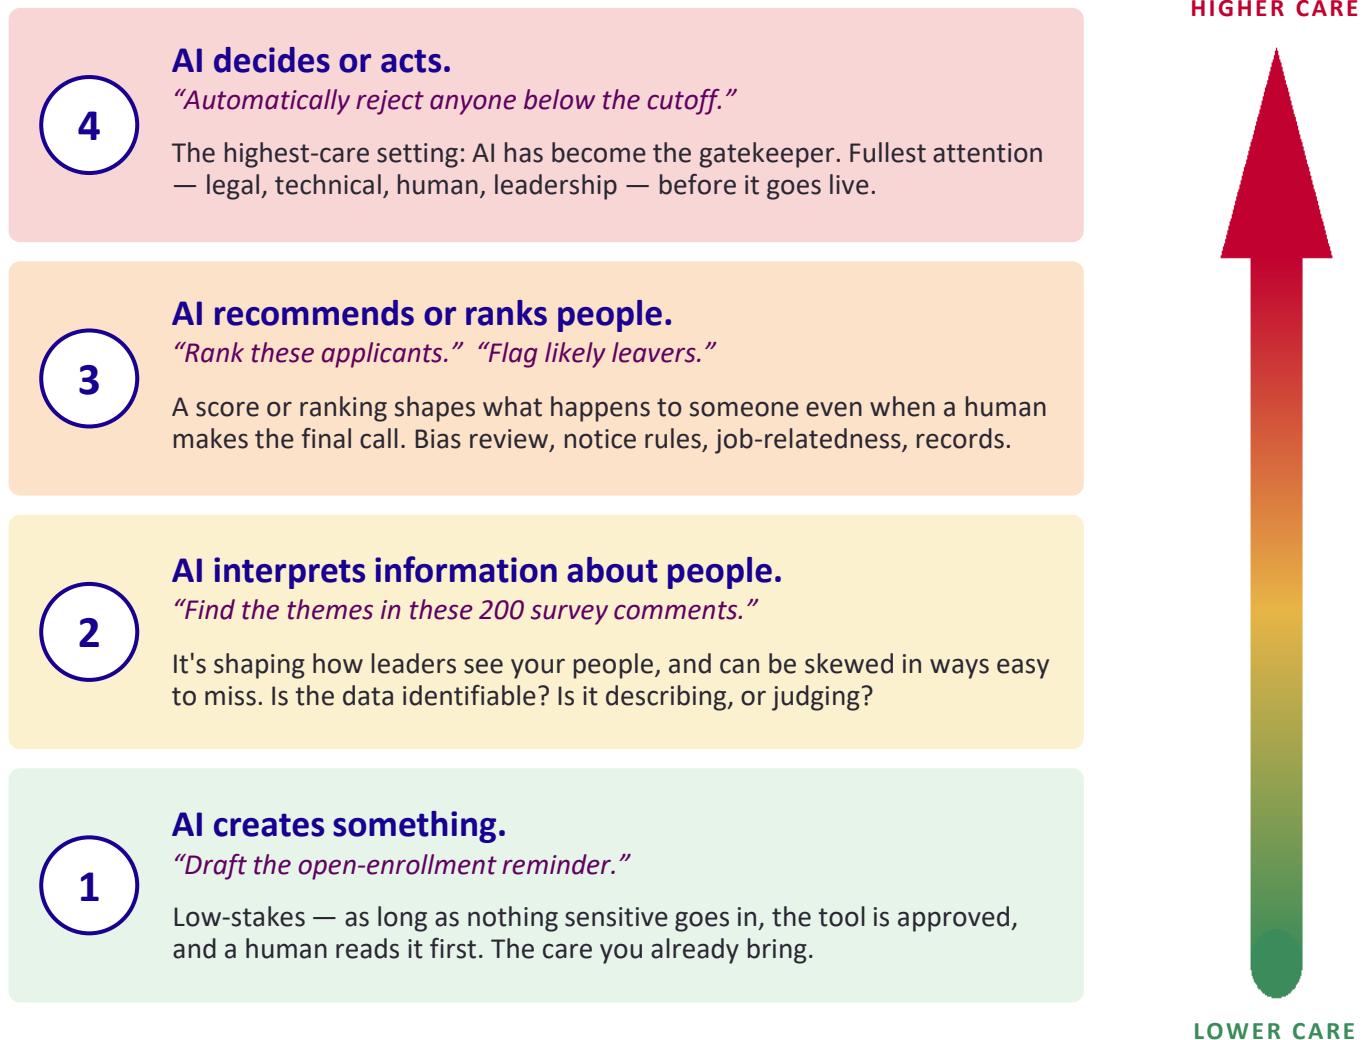
Take Missouri: there's no AI-specific employment law yet. But "my state doesn't have a law" isn't the same as "we're safe." And we all know it's rare for a business to operate — through its people or its vendors — in just one state. So when your applicant is in New York City, your remote employee is in California, or the tool you bought applies one setting to everyone, the law can follow the person, the position, or the activity — not just your business address.

Here's the mindset that keeps you in good standing: run your AI as if these laws already apply to you.

Two reasons. First, it's only a matter of time. AI and its regulations are moving fast, and more states will land here. Second, and bigger: the risk was never only legal. AI sits right where discrimination, bad advice, and compliance all meet, and it can create real exposure whether a statute carries your state's name or not. Operate with integrity, as if you were in the most litigious state on the map. Then you won't be scrambling when the law arrives. You're already there.

How Close Is AI to a Decision?

Think of AI risk as a range, not an on/off switch. As AI moves closer to a real decision about a person, you add a little more care. Read the meter like a climb — the higher the setting, the more care it needs.



AI doesn't have to push the final button to shape the decision.

A score or a ranking moves outcomes too — which is why the newest laws say “materially influence,” not just “fully automated.” And you might already be at Setting 3 or 4 without choosing to be, because the AI doing the ranking is already running inside a system you bought. That's next.

The AI You Didn't Build

Some of the most consequential AI in your organization is already running quietly inside the systems you pay for, and it reaches well beyond HR, into finance, marketing, IT, and operations. As the people who think about both the workforce and the risk, HR is in a great spot to help the whole organization ask good questions about all of it. This is a place HR can lead.

Start with the principle that anchors all of it: outsourced doesn't mean offloaded.

When a system you bought helps make a decision about a person, that decision is still yours — no matter what logo is on the software. Two developments show why that matters right now.

The case that's setting the course —Mobley v. Workday

Whether or not you use Workday, this is the case to watch. Derek Mobley applied to job after job — more than a hundred, by his account — all funneled through employers using Workday's AI screening, and he was turned down every time.



He alleges the system screened people out by race, age, and disability. In May 2025, a federal court let the claim move forward as a collective action under the Age Discrimination in Employment Act (the 1967 law), and the scale is staggering: reporting suggests roughly 1.1 billion applications ran through these tools in the covered window, so the affected group of applicants over 40 could reach into the hundreds of millions. This is potentially one of the largest discrimination collectives ever certified.

Why it matters: the court is allowing the argument that the vendor — not just the employer — could share the liability, because its software does a job employers used to do themselves. “The vendor built it” may not get anyone off the hook, and “we just make the software” may not shield the vendor either.

The bigger idea: algorithmic monoculture

Here's what makes this bigger than one vendor. When lots of employers lean on the same — or similar — AI screening tools, they start making the same call about the same person. A rejection stops being one company's “no” and becomes the market's “no.” Stanford researchers recently found about 10% of applicants to four jobs were rejected by all four — more than chance alone would explain. If a single model is quietly biased and everyone is using it, a qualified person can be locked out everywhere at once.

Which points straight at the real question: what's actually feeding these tools, and how do you keep it in check? That's next.

Ask Before You Trust It

We just saw how one biased model can ripple across an entire market. So before you lean on any tool, here are two things worth understanding about the AI inside it — then the plain questions that get you real answers from any vendor, in HR and beyond.

A word about your data — and about bias.

Most of the time, inside a tool your organization already licenses, your data sits within that program's security — a real benefit of staying in approved tools. Two things are still worth understanding:

Your data may help train their future models. Even inside a secured tool, some vendors use customer data to improve their systems. Worth asking how yours handles it — and where your settings land.

AI is not bias-free. It learns from human data, human data carries human assumptions, and a model can inherit patterns no one intended. So anything involving people deserves human review — and in the big ones (hiring, performance, discipline) your judgment is final, always. That's the thread back to the nondiscrimination laws: the law cares about discriminatory effects, however they arise.

Questions to ask your vendors — HR and beyond.

You don't need to become a software expert — just ask a few plain questions and get clear answers. If a vendor can't answer these clearly, that itself is useful information.

- What AI features are currently active in our account or version?
- Is our data used to train or improve your models — just for us, or platform-wide?
- What decisions is AI making or influencing without a human prompt — and where does human oversight live?
- How is our data handled, protected, and governed inside the platform?

Going deeper, especially for anything touching people decisions:

- Was it validated for our use and workforce?
- Which groups and outcomes were tested?
- Will you tell us before you change the model?
- And when a decision is questioned, what records can you produce?

Find It. Classify It. Control It. Prove It.

If you carry one framework out of today, let it be this one.

Compliance asks what the law requires; governance asks how you decide what AI may do, who approves it, who watches it, and what happens if something goes sideways. A written policy is part of governance — a valuable part — but not the whole thing. Here's a version any HR leader can run.

1 Find it.

Map where AI already lives — well past “who has a ChatGPT account.” Include the models embedded in systems like your ATS and HRIS, plus the smart glasses and notetaking apps people bring in. Ask each vendor which features actually use AI to score, rank, predict, or generate.

2 Classify it.

For each use, **get specific about what it really does**: who it affects, what data it touches, what decision it shapes, and where. Is it creating, interpreting, recommending, or deciding — and if it's wrong, can that outcome be undone?

3 Control it.

Set the ground rules together, before people are off experimenting. Which uses are pre-approved, and which need HR, IT, legal, or privacy to weigh in? What data is allowed in, who reviews the result, who can overrule the tool, and which uses are off the table?

4 Prove it.

Keep the record you'd be glad to have if someone asks: which tool and version, what it was meant to do, the criteria behind a result, the testing that was done, who reviewed it, and when a vendor changed the model underneath you.

The Safe-AI Conversation

Six quick questions that surface whether a use is ready — before it goes live.

- 1 Where is it already used — even in tools we had before a strategy?
- 2 Does it create, interpret, recommend, rank, or decide?
- 3 Could it affect who's hired, paid, promoted, or let go?
- 4 Which people, locations, and jurisdictions does it touch?
- 5 Can we explain its role to a candidate, manager, or regulator?
- 6 Is there real human review — with power to change the result?



A NOTE FROM AI GURU LAUREN PATEL

Governance sounds heavier than it is. Run those four moves on one real use this week and you've already started. You don't need every answer on day one — you need the conversation happening, with HR in the room.

Keep Going

You don't have to memorize the rulebook — just know where the signs are. **Good news: everything here is open to you.** The sample policy is yours to take, and both This Week at Work episodes are free to anyone — no membership required. Here's where to go deeper.

★ Bonus of the day — the AI Sample Policy

Want something concrete to start from? People ask us for this all the time — it's a master template you can shape to your own organization, and it pairs perfectly with the Find/Classify/Control/Prove work.

[Download the AAIM Sample AI Policy →](#)



Members — Train Your Entire Team!

You likely already know about the compliance video series included in your membership through AAIM University. We're pleased to announce we've just added a brand-new training, ***"Responsible AI at Work: Guardrails for Today's Organizations."*** Check out this guided, on-demand training video and set your whole organization on the path to responsible AI usage. All included with your membership.

Contact your advisor to begin today!

This Week at Work

Our free, open-to-everyone series.
Two short episodes that bring today's themes to life:

Smart Glasses at Work: Employer Risks

Governance reaches past ChatGPT and recruiting — consent, biometric data, footage that can be discovered later. [Read the Blog](#) · [Watch the Episode](#)

When AI Becomes Evidence

A court found materials made with a public AI tool weren't protected the way people assumed. A prompt can become a record. [Read the Blog](#) · [Watch the Episode](#)



Your Move Today

Two small moves — plus one for the bold. None of this takes long; it just takes looking.

1. Take a personal AI inventory.

Walk through a normal day and list everything with AI baked in — the meeting notetaker, the scheduling assistant, the smart glasses, your writing tools, your ATS and HRIS. You'll be surprised how long the list gets, and how much of it you never “decided” to adopt.

2. Put one on the meter.

Pick one process you know involves AI and place it on the meter from page 4 — is it creating, interpreting, recommending, or deciding? That one placement tells you how much care it needs, and whether it's getting it.

★ Bonus challenge — start the conversation.

Bring one of these to whoever sets your AI direction — your C-suite, IT, or leadership team. Governance starts the moment someone says it out loud, and HR is the right one to say it.



WHERE THIS GOES NEXT

Getting your organization ready for AI — the inventory, the guardrails, the human-review design, the judgment calls — is a different muscle than using the tools at your own desk. It's exactly the work we help leadership teams think through before they wire AI into how the organization runs. When your team is ready to move from “someone's experimenting with it” to “we've built this to be trusted,” that's the conversation to have. Come find me.