AI in Hiring: You Can Pass the Audit and Still Lose the Lawsuit

What NYC Local Law 144 Now Tells Every Employer Using AI in Hiring

Executive Summary

New York City bet its AI hiring accountability model on one control: the independent bias audit. Three years of evidence show why that control is not enough, and the lesson matters far beyond New York.

Local Law 144, enforced since July 5, 2023, was the first major U.S. law to require independent bias audits for automated employment decision tools used in hiring and promotion. Employers cannot use a covered tool unless it has been audited within the prior year, the audit results are publicly available, and candidates or employees receive required notice.

That was a sensible design. The problem is what happened next.

Three findings now sit on top of each other. A 2024 Cornell-led study of 391 employers found only 18 had posted a bias audit report and 13 had posted a transparency notice. Researchers called the phenomenon "null compliance," because employers retain enough discretion over scope that public silence proves almost nothing. In December 2025, the New York State Comptroller found that the city's enforcement system was ineffective: DCWP reviewed 32 companies and identified one issue of noncompliance, while Comptroller auditors reviewed the same companies and identified at least 17 potential instances. And in *Mobley v. Workday*, a federal court granted preliminary collective certification on an age-discrimination claim involving algorithmic applicant screening, underscoring that AI hiring liability can arise under longstanding employment law whether or not an AI-specific audit regime applies.

Read together, these developments point to one conclusion: the bias audit is a compliance document. It is not a risk control system. Treating it as your AI hiring strategy leaves the organization carrying exposure without enough protection.

This is not a signal to wait for better rules. It is a signal to build an AI hiring governance baseline now: inventory the tools, classify them by decision impact, test outcomes, document human oversight, hold vendors accountable by contract, preserve candidate notice and accommodation pathways, and put high-risk employment automation on the board's agenda.

Govern the hiring decision, not just the tool.

Why This Matters Now

Hiring is one of the most consequential uses of enterprise AI. It determines access to economic opportunity, shapes workforce composition, and sits inside a mature legal framework: Title VII, the ADA, the ADEA, state human rights laws, recordkeeping and validation standards, and privacy obligations.

The adoption curve is steep. SHRM's 2025 Talent Trends research found that 43% of organizations now use AI in HR tasks, up from 26% in 2024. Recruiting is the leading use case among adopters: of organizations that use AI in HR, roughly two-thirds apply it to recruiting and hiring, including job-description drafting, resume screening, candidate search, job-posting customization, and applicant communication.

Most of these tools were bought for speed and consistency. Few were bought with a mature governance model attached.

AI changes the operating reality in three ways.

First, it scales small errors. A poorly calibrated resume screen, keyword model, chatbot, assessment, or interview analyzer can affect thousands of applicants before anyone sees the pattern.

Second, it fragments accountability. Employers rely on vendors for model design, validation, testing, and audit artifacts, but applicants and regulators experience the employer as the decision-maker.

Third, it creates explainability debt. HR leaders can often say a tool is efficient or vendor-validated. Fewer can explain what data it uses, which decision it substantially assists, how a human overrides it, or what adverse-impact testing has shown.

That gap has stopped being only an HR problem. It is now a board problem.

What Local Law 144 Actually Requires

New York City's Department of Consumer and Worker Protection prohibits employers and employment agencies from using an automated employment decision tool unless three conditions are met:

1. The tool has had a bias audit no more than one year before use.
2. A summary of the most recent bias audit is publicly available.
3. Candidates or employees receive required notice that an AEDT will be used, how it will be used, and what data will be collected.

The architecture is public transparency plus independent audit plus candidate notice. That design did something important: it moved AI hiring risk from a private vendor-management question to a public accountability question. Employers could no longer ask only, "Does this tool work?" They also had to ask, "Can we demonstrate that it works fairly, lawfully, and transparently enough to survive scrutiny?"

The design was right. The delivery is where it broke.

The Evidence: A Law Can Require Audits Without Creating Accountability

Start with whether the audits even happen.

In the Cornell-led "Null Compliance" study, 155 trained investigators checked 391 employers as a job seeker would. Eighteen had posted an audit report. Thirteen had posted a transparency notice. Only 11 had posted both in a form that met the law's requirements.

Because employers decide for themselves whether a tool is in scope, a missing audit could mean non-use, a good-faith out-of-scope determination, a report buried where nobody can find it, or noncompliance. The public cannot tell which. Neither can candidates. Often, neither can a board.

A companion study, "Auditing Work," based on interviews with auditors and practitioners, reached the same place: ambiguous definitions, difficulty accessing data, disagreement over the auditor's role, and a transparency-centered theory of change that places too much burden on applicants and outside observers.

Technical critiques add another layer. Researchers analyzing Local Law 144's required metrics have argued that they can miss meaningful distributional differences in candidate scores and therefore under-detect bias in some cases.

None of this means audits are useless. It means audits are necessary and insufficient. A one-time public report cannot substitute for governing the tool across its lifecycle.

The Enforcement Signal: The Regulator Could Not Reliably See Compliance Either

If employers cannot reliably tell what is in scope, can the city tell who is complying? In December 2025, the New York State Comptroller answered: not well enough.

The Comptroller's audit found DCWP's enforcement of Local Law 144 ineffective. DCWP had received only two AEDT complaints during the audit period and had not investigated whether the complaint intake process was working. That process had real gaps: 75% of the auditors' own test complaints placed through the city's 311 system were misrouted and never reached DCWP. DCWP reviewed websites and bias audits for 32 companies and identified one issue of noncompliance. The Comptroller's auditors reviewed the same companies and identified at least 17 potential instances of noncompliance. The audit also found that DCWP officials lacked technical expertise to evaluate AEDT use and did not make adequate use of the city's Office of Technology and Innovation when making AEDT determinations.

This is the predictable arc of AI regulation. When an early law underperforms, regulators rarely abandon the goal. They build capacity.

In May 2026, NYC Council legislation was introduced to create an Office of Artificial Intelligence Oversight inside DCWP, with a complaint portal and authority to recommend enforcement across employment, housing, credit, and services. Phase one was notice, transparency, and self-published audits. Phase two is complaint intake, investigative capacity, and tougher documentation expectations. Phase three is board-level accountability, routine impact assessments, and vendor accountability across jurisdictions.

Weak early enforcement is usually the reason oversight gets stronger, not a reason to relax.

The Liability Signal: The Lawsuit That Does Not Depend on an Audit

Here is why audit-as-strategy is dangerous rather than merely incomplete. While the transparency regime was faltering, one of the most important AI hiring cases in the country advanced on a different theory.

In *Mobley v. Workday*, pending in the Northern District of California, the plaintiff alleges that Workday's algorithm-based applicant screening software produced discriminatory outcomes based on age, race, and disability. On May 16, 2025, the court granted preliminary collective certification on the plaintiff's ADEA age-discrimination claim, finding that the plaintiff had plausibly shown that proposed collective members were subject to a common, unified policy with alleged disparate impact. The potential scale is large: Workday represented in the litigation that roughly 1.1 billion applications were rejected through its tools during the relevant period, putting the possible collective in the hundreds of millions.

Notice what this case does not turn on. Not whether a Local Law 144 audit was posted. Not whether a transparency notice was easy to find. The core litigation theory turns on outcomes: alleged disparate impact under longstanding employment law.

An employer with a current bias audit could still face risk if the underlying hiring system produces discriminatory outcomes, lacks defensible validation, fails to provide accommodation pathways, or cannot explain how human oversight works. The audit and the liability analysis run on separate tracks.

That is the core point. The compliance question and the exposure question are not the same question, and the audit answers only part of the first.

The Broader Pattern: One Decision, Many Overlapping Laws

Local Law 144 sits inside a widening web. Illinois has regulated AI video interviews since 2020 and amended its Human Rights Act in 2024 to cover AI use in employment decisions, with notice obligations effective January 1, 2026. Colorado enacted a broad high-risk AI and algorithmic-discrimination law, though implementation timelines have been debated and revised. California has moved to make explicit that employment-discrimination rules apply to automated decision systems. Federal agencies, including the EEOC and DOJ, have warned that employers can be liable when algorithmic tools create discriminatory barriers, including for applicants with disabilities.

Outside the United States, the EU AI Act classifies employment-related AI as high risk, bringing obligations for risk management, data governance, documentation, transparency, human oversight, accuracy, robustness, and cybersecurity.

There will be no single AI hiring law to comply with. There is a governance stack: employment law, civil rights law, privacy, consumer protection, procurement standards, AI-specific statutes, agency guidance, and litigation. Multistate and multinational employers cannot solve this one jurisdiction at a time.

The Strategic Risk for CHROs and GCs

The most dangerous question an executive can ask is: "Are we using AI in hiring?"

It is too vague to be useful. AI may be embedded in an applicant tracking system, sourcing platform, assessment vendor, chatbot, scheduling tool, background-check workflow, talent marketplace, or job-description tool. It may be marketed as machine learning, automation, matching, scoring, ranking, recommendations, analytics, or optimization.

The better questions are sharper:

1. Which tools influence employment opportunities?
2. Which decisions do they assist: sourcing, screening, assessment, ranking, interview selection, offer, promotion, redeployment, or termination?
3. What data do they use, including inferred or proxy data?
4. Who can override the recommendation, and is override behavior monitored?
5. Which groups may be disadvantaged, including intersectional and disability-related cohorts?
6. What notices, consent flows, and accommodation pathways exist?
7. What evidence shows the tool is job-related, valid, and monitored after deployment?
8. What contractual rights do we have to audit, test, suspend, or obtain data from the vendor?
9. What would we show a regulator, plaintiff's lawyer, journalist, union, board member, or candidate tomorrow?

If those answers are scattered across HR, Legal, Procurement, IT, DEI, and vendors, the organization does not have AI hiring governance. It has fragments.

A Board-Ready Governance Baseline

Employers need a baseline that fits how hiring actually works. Five moves matter most.

1. Build an AI Employment-Tool Inventory

Inventory every tool that touches hiring, promotion, assessment, internal mobility, scheduling, performance, and termination, including anything that only recommends, ranks, prioritizes, summarizes, or screens. Capture owner, vendor, use case, geography, decision stage, data inputs, outputs, human role, affected populations, and contractual rights.

2. Classify Risk by Decision Impact, Not Vendor Label

A "workflow automation" that quietly filters applicants may carry more exposure than a branded AI assistant that only drafts recruiter emails. Classify by consequence: does the tool affect who gets seen, scored, interviewed, advanced, hired, promoted, or removed?

3. Move From Annual Audit to Continuous Evidence

Complete required bias audits for covered tools, then go further. Maintain evidence of validation, adverse-impact monitoring, sample-size limitations, remediation steps, accommodation handling, human review, override patterns, configuration changes, and post-deployment performance. The audit is one artifact in a living control environment.

4. Rebuild Vendor Contracts Around Accountability

AI hiring contracts should include audit rights, data-access rights, change notification, documentation obligations, testing cooperation, incident reporting, retention and deletion terms, accessibility commitments, and clear allocation of responsibility. A vendor that cannot support the employer's legal obligations is not low risk merely because its tool is popular.

5. Put Employment AI on the Board Risk Agenda

Boards do not need model-level detail. They need assurance that management knows where high-impact AI is used, what controls exist, how compliance is tested, what exceptions have surfaced, and what remediation is underway. A quarterly AI employment-risk dashboard should show inventory coverage, high-risk tools, audit status, adverse-impact findings, unresolved vendor gaps, candidate complaints, accommodations, and regulatory developments.

The AIR-L Leadership Lens: From Compliance to Trust

The organizations that handle this well will not be the ones that simply publish the cleanest audit summary. They will be the ones that can explain their hiring system with confidence.

That takes a leadership lens, AIR-L, built on four pillars:

Accountability: A named executive owns employment AI governance, with Legal, HR, Compliance, Procurement, IT, Security, and DEI in defined roles.

Integrity: Tools are tested against the actual job-related criteria they claim to measure, not convenience metrics.

Reliability: Systems are monitored after deployment, especially when labor markets, applicant pools, models, or configurations change.

Legibility: Candidates get meaningful notice, humans understand how to use outputs, and boards get plain-English risk reporting.

Remediation closes the loop across all four. When disparate outcomes or accessibility barriers appear, the organization has a documented path to pause, adjust, replace, or remove the tool. Accountability names the owner. Integrity and reliability generate the evidence. Legibility makes it explainable. Remediation acts on it.

That is the difference between a firm that can defend its hiring system and one that can only produce a certificate.

What to Do in the Next 90 Days

Within 30 days, identify every hiring and promotion tool that may automate, rank, recommend, score, screen, or substantially assist a decision. Stand up a cross-functional owner group led jointly by HR and Legal.

Within 60 days, classify tools by risk, jurisdiction, decision impact, and vendor dependency. For any tool used in NYC hiring or promotion, confirm whether Local Law 144 applies and whether audit, posting, and notice obligations are met.

Within 90 days, brief the executive team or board risk committee on the inventory, top exposures, audit gaps, vendor-contract gaps, candidate-notice practices, accommodation pathways, and remediation plan.

Do not wait for a complaint, or a collective action, to learn how your hiring technology actually works.

Conclusion

Local Law 144 is not the final form of AI hiring regulation. It is the first serious draft, and its weaknesses are exactly what make it instructive. Employers struggled to determine what was in scope. The city struggled to detect compliance. And while everyone debated audits, liability continued to move under existing employment law.

The takeaway is not "audit harder." The takeaway is that the bias audit is the starting line, not the control system.

Employers that build governance around the hiring decision, including inventory, impact classification, continuous evidence, vendor accountability, and board oversight, will be better positioned to use AI responsibly, defend their practices, keep hiring fast, and preserve trust with candidates, regulators, and the public.

Those who keep mistaking the audit for the strategy will keep carrying the risk without enough protection.

Suggested LinkedIn Post Intro

The bias audit was supposed to be the safeguard for AI in hiring. The evidence now says it is only one piece of the control system.

NYC's Local Law 144 made bias audits mandatory for covered automated hiring tools. Three years in, a Cornell-led study found very few employers posting required audit and transparency materials, the New York State Comptroller found the city's enforcement system ineffective, and one of the country's most important AI hiring lawsuits is advancing under traditional employment-discrimination law.

For CHROs and General Counsel, the question is no longer simply, "Has the tool been audited?"

The better question is whether the organization can explain, test, monitor, and defend the entire hiring system around it.

We put together a practical executive brief on what Local Law 144 really signals, and why the next move is a Board AI Governance Baseline.

Sources and Further Reading

- NYC Department of Consumer and Worker Protection, "Automated Employment Decision Tools (AEDT)": https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page
- NYC Rules, "Automated Employment Decision Tools (Updated)," adopted rule effective July 5, 2023: https://rules.cityofnewyork.us/rule/automated-employment-decision-tools-updated/
- Office of the New York State Comptroller, "Enforcement of Local Law 144 - Automated Employment Decision Tools," Dec. 2, 2025: https://www.osc.ny.gov/state-agencies/audits/2025/12/02/enforcement-local-law-144-automated-employment-decision-tools
- Wright, Muenster, Vecchione, Qu, Cai, Matias et al., "Null Compliance: NYC Local Law 144 and the Challenges of Algorithm Accountability," FAccT 2024: https://arxiv.org/abs/2406.01399
- Groves, Metcalf, Kennedy, Vecchione, Strait, "Auditing Work: Exploring the New York City algorithmic bias audit regime," 2024: https://arxiv.org/abs/2402.08101
- Hilliard, Koshiyama et al., "Local Law 144: A Critical Analysis of Regression Metrics," 2023: https://arxiv.org/abs/2302.04119
- Mobley v. Workday, Inc., No. 3:23-cv-00770 (N.D. Cal.), Civil Rights Litigation Clearinghouse: https://clearinghouse.net/case/44074/
- SHRM, "The Role of AI in HR Continues to Expand" (2025 Talent Trends): https://www.shrm.org/topics-tools/research/2025-talent-trends/ai-in-hr
- New York City Council, Int. 0919-2026, Office of Artificial Intelligence Oversight: https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=8034442&GUID=669B5731-411B-40A9-99F2-6AD05E945AE8