NOVARA
CONSULTING GROUP

# The Development Exception: Internal Mobility Platforms and the Perimeter of Employment AI Governance

Heather Grizzle · May 21, 2025

*Internal mobility platforms rank employees for promotion, which puts them inside NYC Local Law 144 and EU AI Act Annex III. Most employers assume otherwise.*

**Contents**

1. The boundary no statute draws
2. What the perimeter actually covers
3. The evidentiary problem
4. The bias claim requires evidence, not assurance
5. Governance implications
6. References

## The boundary no statute draws

Institutional AI governance in employment has converged, with notable consistency, on a boundary that appears nowhere in the governing law. The boundary separates hiring from development. Systems that evaluate external applicants are treated as decision technology and receive the corresponding apparatus: bias auditing, candidate notice, vendor documentation, and legal review. Systems that evaluate existing employees for internal opportunity are treated as development technology and are procured under review standards appropriate to a learning management platform.

The distinction is administratively intuitive. External candidates are strangers to the organization, and the organization owes them process. Internal candidates are already employed, already known, and already inside the relationship. Extending opportunity to them presents as a benefit conferred rather than a decision rendered. The distinction is also organizationally reinforced, because internal mobility platforms are typi-

cally sponsored by talent development functions, budgeted against engagement and retention objectives, and evaluated on adoption metrics rather than decision quality.

Neither the technology nor the statutory language supports the distinction. An internal mobility platform that infers employee capability from resumes, project histories, learning records, performance data, and behavioral signals, and then ranks employees against internal openings, is generating an assessment that materially influences promotion, compensation, and assignment. That is a consequential decision about a person. The circumstance that the person volunteered for the platform, and that the organization describes the output as a recommendation, does not alter the character of the determination being made.

## What the perimeter actually covers

New York City's Local Law 144 has been enforceable since July 2023 and reaches automated employment decision tools used to substantially assist or replace discretionary decision-making for employment decisions. The statute is explicit that it covers hiring **and promotion**. Its obligations are procedural rather than substantive: an independent bias audit conducted within the prior year, public posting of the audit summary, and advance notice to affected individuals. Penalties accrue per violation, and continued use is treated as a continuing violation.

An organization that has commissioned bias audits for its external applicant screening while treating its internal talent marketplace as out of scope has drawn a line the ordinance does not recognize. Where the internal system scores or ranks employees for postings that constitute promotion, the audit obligation attaches. Public evidence does not yet indicate how consistently enforcement authorities have reached internal mobility deployments, but the absence of enforcement to date is a poor foundation for a compliance position, particularly where the exposure compounds daily.

The pattern extends beyond New York. Illinois HB 3773, enacted in August 2024 and scheduled to take effect January 1, 2026, amends the Illinois Human Rights Act to prohibit the use of artificial intelligence that produces discriminatory effects on protected classes, to bar the use of zip codes as proxies for protected characteristics, and to require notice whenever artificial intelligence is used in decisions concerning recruitment, hiring, promotion, discipline, or discharge. The notice obligation attaches to use, not to outcome. The discrimination provision is enforceable through the Human Rights Act, which carries a private route to relief.





Colorado's SB 24-205, enacted in May 2024 and scheduled to take effect February 1, 2026, identifies employment decisions as consequential decisions and imposes on deployers a duty of reasonable care against algorithmic discrimination, together with risk management programs, impact assessments, and disclosure obligations. The statute was signed with an accompanying letter from the governor urging legislative revision before implementation, and it has attracted sustained amendment pressure. Its final form is uncertain. That uncertainty is itself a governance consideration rather than a reason for deferral, because an organization that calibrates its internal controls to whichever regime proves weakest will be recalibrating them at each legislative session.

California's Civil Rights Council has been developing regulations addressing automated decision systems in employment, and the direction of that rulemaking is consistent with the pattern above. In the European Union, Regulation (EU) 2024/1689 entered into force in August 2024 and designates artificial intelligence systems intended for use in employment, worker management, and access to self-employment as high risk under Annex III, expressly including systems used to make decisions affecting promotion, task allocation, and the evaluation of work performance. Obligations for those systems are scheduled to apply from August 2026. The Annex III language is directed at the function performed rather than the product category marketed, and internal mobility engines perform the enumerated functions.

July 2023
**NYC Local Law 144 enforceable (hiring and promotion)**

August 2024
**Regulation (EU) 2024/1689 enters into force**

1 January 2026
**Illinois HB 3773 takes effect**

1 February 2026
**Colorado SB 24-205 takes effect**

August 2026
**EU AI Act high-risk obligations apply**

## The evidentiary problem

The business case for AI-enabled internal mobility circulates in a form that will not survive procurement scrutiny. Figures asserting that these platforms reduce attrition by a third, or that engagement tools will lift retention by a quarter within a stated year, propagate through vendor content aggregations in which each claim carries a source attribution that resolves to another aggregation. The underlying study, where one is sought, generally does not exist.

A second and more consequential distortion involves the misattribution of genuine research. LinkedIn′s Workplace Learning Report and its associated Career Development Index find that organizations scoring in the top quintile on career development maturity, measured through stated commitments, job posting language, leadership skill prevalence, and internal mobility levels, show materially higher rates of promotion into leadership, longer employee tenure, and greater learner engagement than organizations in the bottom quintile. That is a finding about organizational investment in career development. It is routinely restated in vendor materials as a finding about the effect of artificial intelligence tooling, which is not what was measured and not what the comparison groups distinguish.

What the credible record supports is narrower. Several large enterprises, including Unilever, HSBC, and Schneider Electric, have publicly attributed meaningful reductions in voluntary attrition to internal talent marketplace deployments. Analyst forecasts have projected substantial adoption among large enterprises. Case evidence indicates faster internal fill times and reduced contractor spend. These are self-reported operational outcomes and forecast estimates rather than controlled evaluation, and they should be characterized as such in any business case that will be read by an auditor, a regulator, or opposing counsel. The distinction between an outcome an organization observed after deployment and an effect attributable to the system is the distinction that governance exists to preserve.

## The bias claim requires evidence, not assurance

The most common justification advanced for algorithmic internal mobility is that it reduces bias. The argument has a defensible premise. Discretionary manager-controlled promotion carries well-documented disparities, talent hoarding suppresses employee

visibility, and opacity about internal openings distributes opportunity unevenly. Transparency about what roles exist and who may apply has independent value.

Consistency is not equity.

The conclusion does not follow from the premise. An internal mobility system trained on an organization's own advancement history will reproduce the patterns in that history with greater consistency, broader application, and less traceability than the managers whose judgment it supplements. Consistency is not equity. The scrapping of Amazon's experimental recruiting tool, after it was found to disadvantage women, remains the standing demonstration that historical training data encodes historical outcomes, and the case involved external hiring, where the training signal is comparatively sparse. Internal mobility systems operate on the accumulated record of the employment relationship itself, including performance ratings, project assignments, and manager assessments, each of which carries its own distributional history.

The governance consequence is that a claim of bias reduction is a factual assertion requiring evidence, not a design intention that may be recited. Under a disparate impact standard of the kind Illinois has adopted, an employer's good-faith belief that a tool improves fairness is not responsive to the question of whether outcomes differ across protected classes. Only outcome testing is responsive to that question.

## Governance implications

The operative recommendation is that internal mobility systems be brought inside the employment AI governance perimeter at procurement rather than at incident.

Procurement documentation should require the supplier to state, in contractual language, whether system outputs are used to make, guide, or materially influence decisions concerning promotion, compensation, assignment, or continued employment. That formulation, rather than the product category under which the system is marketed, determines regulatory coverage across every regime surveyed above. A supplier unwilling to characterize its own outputs in those terms has communicated something useful.

Bias audit documentation should be required as a condition of award rather than requested after deployment, and should identify the audit methodology, the protected

class categories examined, whether intersectional analysis was performed, the population on which the audit was conducted, and the date. An audit conducted on a vendor's aggregate client population is not an audit of the deploying organization's outcomes, and the distinction should be recorded rather than assumed away.

Model documentation should identify training data provenance and the features on which inferences are drawn, with explicit confirmation that no geographic, educational, or tenure-derived variable operates as a proxy for a protected characteristic. Notice and adverse action workflows should be specified contractually rather than left to implementation, because the notice obligations under the statutes surveyed attach to the deploying employer regardless of what the vendor has built. Indemnification language warrants particular attention, as the emerging pattern in state legislation constrains the transfer of discrimination liability to suppliers.

☐ Supplier states contractually whether outputs make, guide, or materially influence decisions on promotion, compensation, assignment, or continued employment

☐ Bias audit documentation required as a condition of award, with methodology, protected class categories, intersectional analysis, population, and date

☐ Model documentation identifying training data provenance and inference features, with confirmation that no geographic, educational, or tenure-derived variable acts as a proxy

☐ Notice and adverse action workflows specified contractually

☐ Indemnification language reviewed against state limits on transferring discrimination liability

None of this constitutes an argument against AI-enabled internal mobility. The talent rationale is substantive, and the alternative of opaque discretionary gatekeeping is not evidently superior. The argument is that a category of system marketed as an employee experience initiative is functionally consequential decision infrastructure, that the statutory perimeter already reaches it, and that organizations discovering this after deployment will be retrofitting governance onto tools that have been shaping advancement for several cycles. The least expensive point at which to govern such a system is before it is acquired.

## References


- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence, Annex III(4).
- New York City Administrative Code § 20-870 et seq. (Local Law 144 of 2021); Department of Consumer and Worker Protection final rules, effective 5 July 2023.
- Illinois House Bill 3773, 103rd General Assembly (2024), amending the Illinois Human Rights Act, effective 1 January 2026.
- Colorado Senate Bill 24-205, Concerning Consumer Protections in Interactions with Artificial Intelligence Systems (2024), effective 1 February 2026.
- LinkedIn Learning, *Workplace Learning Report 2025*, and the associated Career Development Index methodology.
- Dastin, J. "Amazon scraps secret AI recruiting tool that showed bias against women." Reuters, 10 October 2018.
- McKinsey & Company, "Stave off attrition with an internal talent marketplace," McKinsey Talks Talent.


**Disclaimer.**
Novara Consulting Group provides governance and procurement advisory services for institutional artificial intelligence deployment. This article is analytical and does not constitute legal advice.

**TIPS**

Know something about a sign language AI system that the public should know? Send it to tips@novaracg.com. Every tip is read by me and kept confidential. If you work in the UK, writing to us is not by itself a protected disclosure under the Public Interest Disclosure Act 1998, so before you share employer information, check your options with Protect (protect-advice.org.uk), the free whistleblowing advice line.

## How to cite

Grizzle, H. M. (2025, May 21). The Development Exception: Internal Mobility Platforms and the Perimeter of Employment AI Governance. Novara Consulting Group. https://www.novaracg.com/2025/05/21/the-development-exception-internal-mobility-platforms-and-the-perimeter-of-employment-ai-governance/