

This document was translated from English by machine and may contain mistakes. The English original is at <https://www.novaracg.com/2025/07/03/the-queue-exists-for-a-reason-agentic-ai-and-no-code-deployment-in-human-resource-operations/>.

The Queue Exists for a Reason: Agentic AI and No-Code Deployment in Human Resource Operations

Heather Grizzle · July 3, 2025

An AI agent that screens resumes with no human reviewing exclusions is the paradigm automated employment decision tool, not an edge case. Local Law 144 applies. (159) No-code agent deployment removes the IT queue. The queue was where data assessment, vendor review, and retention specification happened. Those obligations remain. (155) Text sentiment analysis is not banned under the EU AI Act. Voice and video pulse checks are. The line moves with modality, not function. (137) Does the agent assess people or serve them? One question separates onboarding automation from a regulated employment decision system. (135) Your agents log every exclusion they make. That record is discoverable, and whether it helps or hurts was decided at configuration by someone not thinking about it. (159)

目录

1. The capability claim is largely sound
2. What autonomy changes about legal position
3. Where the sentiment agent actually sits
4. The agent that is genuinely fine
5. Governance implications
6. 参考文献

The capability claim is largely sound

The operational case for agentic systems in human resource work does not require much defence. A system that executes a multi-step workflow without a human initiating each step is meaningfully different from a rules engine, and the difference is not rhetorical. Onboarding sequences, document distribution, scheduling, status



communication, and access provisioning are genuinely well suited to autonomous execution, and organizations that automate them recover capacity that was never usefully spent.

The no-code development environment is likewise a real change rather than a marketing one. Human resource professionals understand the employee lifecycle in ways that engineering teams building against a requirements document do not, and configuration tools that put design in the hands of the people who understand the process will produce better processes.

Neither observation is in dispute here. What is in dispute is the inference commonly drawn from them, which is that because human resource functions can now deploy autonomous systems without technical intermediation, they should. That inference treats the technical intermediary as a bottleneck. It is more accurately understood as a control, and organizations that remove it without replacing what it did have not accelerated their operations. They have relocated a set of decisions to a function that has not been resourced to make them.

What autonomy changes about legal position

The distinction that matters is not between automation and agency. It is between systems whose outputs inform a decision about a person and systems whose outputs do not, and agentic architectures tend to move systems across that line without anyone noticing the crossing.

Consider the top-of-funnel recruiting agent that has become the standard illustration of the category. It filters applications against stated qualifications, conducts pre-screening, communicates status to candidates, and schedules interviews with those it advances. The benefit is typically framed as follows: hiring managers now interact only with pre-vetted candidates. That framing describes the compliance exposure precisely, though not intentionally. If hiring managers see only survivors, then no person has reviewed the excluded population, and the exclusion is wholly attributable to the system.

New York City's Local Law 144 reaches automated employment decision tools used to substantially assist or replace discretionary decision-making, and requires an

If hiring managers see only survivors, then no



independent bias audit conducted within the preceding year, public posting of the audit summary, and advance notice to candidates.

A screening agent operating without human review of exclusions is not at the margin of that definition. It is the paradigm case.

Illinois HB 3773, effective January 2026, will require notice whenever artificial intelligence is used in recruitment or hiring decisions and will make discriminatory effects actionable

under the Human Rights Act, with the notice obligation attaching to use rather than to outcome. Colorado's statute, presently scheduled for early 2026 though subject to active amendment pressure, identifies employment decisions as consequential decisions and imposes employer obligations. Annex III of the EU Artificial Intelligence Act designates employment systems as high risk, with obligations scheduled from August 2026.

None of these regimes contain an exception for systems the organization built itself. The obligations attach to the deployer. An organization that acquires a screening tool from a vendor at least receives a contract, a vendor security review, and in many cases audit documentation the vendor produced for other clients. An organization that assembles the same functionality from a no-code builder receives none of those artifacts and holds the same obligations.

This is the specific sense in which the queue was a control. Technical intake processes are where data protection assessments are triggered, where vendor risk review occurs, where retention and export requirements are specified, and where someone asks what the system does with personal data. Those functions were not features of the queue. They were its purpose. Removing the queue removes them, and the obligations they existed to discharge remain.

Where the sentiment agent actually sits

The engagement monitoring agent, which administers pulse surveys, analyses response sentiment, and alerts leadership when morale declines within a department, warrants more careful treatment than it usually receives, because the risk is commonly misidentified.

person has reviewed
the excluded
population, and the
exclusion is wholly
attributable to the
system.



It is worth stating plainly that text-based sentiment analysis of survey responses is not prohibited under the European Union's ban on workplace emotion recognition. The prohibition in Article 5(1)(f) is tied to the Act's definition of an emotion recognition system, which requires inference from biometric data, and the Commission's guidelines on prohibited practices state that a system inferring emotions from written text is not based on biometric data and falls outside the prohibition's scope. Organizations advised otherwise have been advised imprecisely.

MODALITY, NOT FUNCTION

The scope changes with modality rather than with function: a written pulse survey falls outside the prohibition, while voice, recorded meetings, or video inputs bring it inside.

The scope changes with modality rather than with function. If the pulse check is administered by voice, or if the engagement signal is drawn from recorded meetings, video interviews, or any input from which vocal or facial characteristics can be analysed, the analysis is inference from biometric data and the prohibition applies. It applies as a prohibition rather than as a high-risk classification, which means it is not remediable through documentation. Organizations extending sentiment capability from written surveys into meeting analysis are crossing a line they may not know exists, and the extension is exactly the kind of incremental scope expansion that agentic deployment makes easy.

The exposure that does attach to the text-based version is domestic and distinct. Two features of the design create it. The first is de-anonymization through granularity: an alert stating that morale has declined within a named department is, in a department of modest size, functionally identifiable, and it is delivered to the leadership of that department. Engagement instruments depend on a confidentiality representation, and a system that routes sentiment findings to the managers whose conduct is being assessed has withdrawn that representation without disclosing the withdrawal. The second is that automated monitoring which detects and escalates collective employee dissatisfaction operates in the territory that Section 7 of the National Labor Relations Act protects. Surveillance of protected concerted activity is an unfair labor practice, and the statutory principle does not depend on the enforcement posture of any particular General Counsel, which has shifted and will shift again. An organization



that has built a system to detect emergent collective grievance and route it to management has built something it should have counsel look at before it runs.

The agent that is genuinely fine

Not every agentic deployment carries this profile, and the distinction is worth drawing because it is the distinction organizations should be applying.

The onboarding agent that distributes materials, assigns and tracks training completion, and provisions system access after an offer is accepted makes no assessment of any person. The employment decision has already been made. The agent executes logistics consequent to it. There is no ranking, no scoring, no exclusion, and no inference about the individual. It sits outside every regime discussed above, and an organization that deploys it has automated administration rather than judgment.

The lesson is not that agentic systems are dangerous. It is that the category contains systems with entirely different regulatory profiles, and that the profile is determined by whether the system assesses people or serves them. That question is answerable in a sentence per system, by someone who knows to ask it, at the moment of design. It becomes considerably harder to answer eighteen months later across a portfolio of agents that individual teams built and nobody catalogued.

Governance implications

The operative recommendation is that no-code agent deployment requires a lightweight intake control specifically because it has removed the heavyweight one, and that the control should be designed to be fast enough that it is not itself routed around.

A single screening question is sufficient for the majority of cases. Does this agent's output influence a decision about a person, including selection, exclusion, ranking, scheduling priority, escalation, or assessment? Agents answering no proceed. Agents answering yes enter a defined path: coverage determination against applicable jurisdictions, notice requirements identified, audit obligation established, human review point specified, and an owner recorded. The path should be measured in days rather than weeks, because a control that reinstates the delay it replaced will be evaded exactly as the original was.



Screening question

Does this agent's output influence a decision about a person, including selection, exclusion, ranking, scheduling priority, escalation, or assessment?

→ **Decision about a person?**

Decision about a person?

Yes → Defined path No → Proceed

Proceed

Agents answering no proceed.

Defined path

Coverage determination against applicable jurisdictions, notice requirements identified, audit obligation established, human review point specified, and an owner recorded.

Organizations should also maintain a register of deployed agents, for the unglamorous reason that they will otherwise not have one. The distributed configurability that makes these platforms valuable also means no central function holds an inventory, and an organization that cannot enumerate its automated decision systems cannot determine its obligations under any of the regimes described above. The register is not a governance artefact in the abstract. It is the precondition for answering a regulator's first question.

Finally, organizations should attend to what these systems record. Agentic platforms log their own execution, and those logs are discoverable. An agent that excluded four hundred applicants has created a detailed record of the basis on which it did so, held by the organization, produced by the organization, and available to anyone who asks for it in litigation. That record is either the organization's best evidence or its worst, and which one it is was determined at configuration by someone who was not thinking about it.



The premise that human resource functions should control their own technology is correct. Control includes the obligations that come with it.

参考文献

Regulation (EU) 2024/1689, Article 3(39), Article 5(1)(f), and Annex III(4); European Commission, *Guidelines on Prohibited Artificial Intelligence Practices*, February 2025.

New York City Administrative Code § 20-870 et seq. (Local Law 144 of 2021).

Illinois House Bill 3773, 103rd General Assembly (2024), amending the Illinois Human Rights Act, effective 1 January 2026.

Colorado Senate Bill 24-205 (2024).

National Labor Relations Act § 7, 29 U.S.C. § 157; § 8(a)(1), 29 U.S.C. § 158(a)(1).

Novara Consulting Group, "The Development Exception: Internal Mobility Platforms and the Perimeter of Employment AI Governance," May 2025.

* * *

Disclaimer.

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

引用方式

Grizzle, H. M. (2025, July 3). The Queue Exists for a Reason: Agentic AI and No-Code Deployment in Human Resource Operations. Novara Consulting Group. <https://www.novaracg.com/zh/2025/07/03/the-queue-exists-for-a-reason-agentic-ai-and-no-code-deployment-in-human-resource-operations/>

