

The Policy Is the Exhibit: Generative Artificial Intelligence in Human Resources Documentation and Advice

Heather Grizzle · May 27, 2025

Generative AI drafts handbooks, policies, and job descriptions fast. Those documents are evidence in litigation, and provenance is what gives them weight.

Contents

1. Where the exposure actually sits
2. The policy is the exhibit
3. Hallucinated authority, and why human resources is worse placed than legal
4. Continuous alignment is not how the technology works
5. Audit readiness is inverted
6. The job description is a legal instrument
7. The chatbot gives advice the employer is bound by
8. What to establish
9. Conclusion
10. References

Where the exposure actually sits

Most attention to artificial intelligence in human resources has concentrated on selection: who gets screened, ranked, promoted, and on what basis. That focus is warranted and this firm has written at length about it. It has also left the fastest-growing application almost

entirely unexamined.

Generative systems are now being used to draft, update, and localize the documents that constitute an employer's legal position. Handbooks. Accommodation procedures. Leave policies. Job descriptions. Disciplinary standards. Investigation templates. Multi-state compliance addenda. The pitch is that this work, historically slow and expensive, can be produced instantly and kept continuously current as regulations change.

Novara Consulting Group's position is that this application carries greater and more immediate risk than automated screening, for a reason specific to the artifact: a policy is not an operational document. It is evidence. It is produced in litigation, cited in agency investigations, relied upon by employees, and read against the employer by people whose job is to find the gap between what it says and what happened. An organization that has generated its legal position without provenance has created a large volume of admissible material about itself that nobody at the organization can account for.

The policy is the exhibit

In an employment dispute, the employer's written policy performs several functions at once, and each of them cuts.

It establishes what the employer told employees, which matters for reliance and, in many states and depending on drafting and disclaimers, for contract formation. It establishes the standard the employer set for itself, against which its conduct in the specific case will be measured; an employer that failed to follow its own stated procedure is in a materially worse position than one whose procedure was silent. It establishes what the employer knew, since a policy addressing a legal requirement demonstrates awareness of it. And it establishes consistency or its absence, because a policy applied differently to different employees is the ordinary evidentiary route to a discrimination finding.

None of that is new. What is new is that the document may now have been produced by a system that optimizes for plausibility, by a person who is not a lawyer, in a process that leaves no record of what authority was consulted or what judgment was exercised.

When opposing counsel asks who drafted the accommodation procedure, on what basis a particular provision was included, and what legal authority supports it, the answer that a generative system produced it and someone

reviewed it for tone is not a good answer. It is not merely unhelpful; it establishes that the employer adopted a legal position without determining whether it was correct.

It establishes that the employer adopted a legal position without determining whether it was correct.

Hallucinated authority, and why human resources is worse placed than legal

The tendency of generative systems to produce confident and entirely fabricated legal citations is documented and consequential. In 2023 a federal court in the Southern District of New York sanctioned attorneys who filed a brief containing judicial decisions that did not exist, generated by a chatbot and unverified. Comparable incidents have followed in numerous jurisdictions, and the pattern is now familiar to anyone practicing.

The important observation for this discussion is that those were lawyers. They had the training to recognize that a citation must be checked, access to the databases required to check it, and professional obligations that made failure sanctionable. The error still occurred repeatedly.

A human resources generalist drafting a multi-state leave policy has none of those advantages. She frequently cannot tell whether a cited provision exists, does not have the research tools to confirm it, and has no professional duty of verification attaching to her

work. The output arrives fluent, structured, appropriately hedged, and formatted like every competent policy she has read. Fluency is precisely what these systems produce best and precisely what a non-specialist reader uses as a proxy for accuracy.

The errors that matter most are also the least visible. A fabricated citation is at least checkable once someone thinks to check. A subtly misstated standard is not: a policy that describes the interactive process as beginning when an employee submits documentation, rather than when the employer becomes aware of the need, reads perfectly and is wrong in a way that will produce a failure-to-accommodate finding two years later.

Continuous alignment is not how the technology works

The claim that a generative system adapts to changing regulations across states and countries, ensuring continuous alignment, should be treated as false unless the vendor can describe a specific mechanism.

A language model has a training cutoff. Absent retrieval against an authoritative and maintained legal source, it produces the law as it appeared in training data, with no signal to the user that the answer is stale. Employment law is among the worst possible domains for this limitation: pay transparency requirements, paid leave entitlements, non-compete restrictions, salary history bans, and scheduling rules have proliferated and changed across jurisdictions at a pace that outruns any static corpus. A policy generated today may describe a requirement that was amended eighteen months ago, and it will describe it confidently.

The multi-jurisdiction case is the worst one rather than the best one, contrary to how it is marketed. Localization requires knowing which of fifty-plus regimes applies, how they interact, and where the local requirement exceeds the federal floor. It is the task most dependent on current specialist knowledge and the one where an error is least likely to be noticed by the person requesting the output, because she asked for the localization

precisely because she did not know.

An organization should also recognize the claim's shape. In January 2025 the Federal Trade Commission required an accessibility software provider to pay one million dollars over representations that its automated product could make any website conform to accessibility standards and maintain that conformance as content changed. Continuous automated compliance with a technical standard is the same claim, made about a different standard, and it did not survive contact with a regulator.

Audit readiness is inverted

The proposition holds that generated documentation produces audit readiness. It produces the opposite, and understanding why clarifies what these engagements actually test.

An auditor, an agency investigator, or an opposing expert does not primarily ask whether a policy exists. Existence is trivially demonstrated. The questions are how the organization determined that this policy satisfies the applicable requirement, who made that determination, what authority they relied on, when it was last reviewed against current law, and what changed at the last revision and why.

Generated documentation answers none of these. It produces the artifact while destroying the record that gives the artifact weight. An organization with a thin, dated, plainly drafted policy and a documented annual review by named counsel is in a substantially stronger position than one with a comprehensive, current-looking, fifty-state policy set whose origin nobody can describe.

THE REMEDY

The remedy is not to avoid the tools. It is to maintain provenance: what was generated, what was verified, by whom, against what source, on what date. That record costs almost nothing to keep at the time of drafting and cannot be reconstructed afterward.

The job description is a legal instrument

One item in the standard list deserves separate treatment, because it is routinely treated as marketing copy and is in fact one of the most consequential documents the function produces.

Under the ADA (Americans with Disabilities Act), a qualified individual is one who can perform the essential functions of the position, and the implementing regulations identify written job descriptions prepared before advertising or interviewing as evidence of what those essential functions are. A generated job description that lists duties nobody actually performs, or that includes physical or sensory requirements copied from a generic template, does two things simultaneously: it screens out candidates who could perform the job, and it supplies the plaintiff with the employer's own document showing that the stated requirement was never essential.

The same document carries weight in exempt classification analysis under the Fair Labor Standards Act, where the inquiry turns on actual duties but the written description is evidence of what the employer represented those duties to be. It is also the source against which any selection procedure's job-relatedness is assessed under the Uniform Guidelines.

Generated job descriptions are attractive because writing them is tedious and because generated ones read well. They are also the point at which a small drafting convenience

becomes an evidentiary problem, and the physical-requirements section is where this happens most often.

The chatbot gives advice the employer is bound by

The final application, an artificial intelligence assistant answering employee policy questions on demand, transfers a specific and underappreciated risk.

An employee who asks about leave eligibility, accommodation procedure, or complaint routing and receives an answer has received it from her employer. If the system tells her she is not eligible for leave she is entitled to, and she does not take it, the employer has interfered with a statutory right, and interference under the Family and Medical Leave Act does not require intent. If it tells her the accommodation process requires medical documentation she is not obliged to provide, and she abandons the request, the employer has a failure-to-accommodate problem with a written record of the deterrent. If it directs a harassment complaint into a channel that never reaches a decision-maker, the employer's notice defense is compromised by its own system.

Three controls follow directly. Conversations should be logged and retained, because they are employer statements and will be sought. Any query touching leave, accommodation, complaint, or investigation should route to a named human rather than being answered, since these are the categories where an incorrect answer is actionable and where the value of an instant response is lowest. And the interface should state plainly that it provides general information and does not determine eligibility or entitlement.

What to establish

Six items, before a generative system produces anything the organization will rely on or an employee will read.

1. Keep a provenance record for every generated document: what was produced, what source was checked, by whom, on what date, and what was changed on review. Undated and unattributed policy is the failure mode.
2. Require legal review for anything stating a legal standard, entitlement, or procedure, and treat the reviewer's name and date as part of the document rather than as an internal note.
3. Ask any vendor claiming regulatory currency to describe the mechanism: what authoritative source is retrieved, how often, and how the user is alerted when a provision has changed. If the answer is that the model was trained on legal material, the product does not do what is claimed.
4. Treat job descriptions as legal instruments. Review the essential functions and any physical, sensory, or communication requirements against what incumbents actually do, and date the review.
5. Log and retain employee-facing assistant conversations, and route leave, accommodation, complaint, and investigation queries to a person.
6. Do not represent generated documentation as audit readiness, internally or externally. Audit readiness is the review record, not the document.

Conclusion

The case for generative assistance in human resources documentation is real. Drafting is genuinely slow, the function is genuinely under-resourced, and a well-supervised first draft is a legitimate use of the technology.

What makes this application different from the others is that its output is evidence.

Screening systems produce decisions the employer may struggle to reconstruct.

Documentation systems produce statements the employer will be held to, in its own voice, in writing, in volume, at speed. The organization that generates its handbook has not

automated a chore. It has authored its side of a future dispute without reading what it said, and it will encounter that document again under conditions it does not control.

The discipline required is modest: verify, attribute, date, and keep the record of who decided. That discipline is the entire difference between using the tool and being exposed by it.

References

Americans with Disabilities Act of 1990, 42 U.S.C. § 12111(8); 29 C.F.R. § 1630.2(n)(3), evidence of essential functions including written job descriptions prepared before advertising or interviewing.

Family and Medical Leave Act, 29 U.S.C. § 2615(a)(1), prohibition on interference with protected rights.

Fair Labor Standards Act, 29 U.S.C. § 201 et seq., exemption analysis based on actual duties.

Federal Trade Commission, *In the Matter of accessiBe, Inc.*, File No. 2223156, complaint and proposed order announced January 3, 2025.

Mata v. Avianca, Inc., No. 22-cv-1461 (S.D.N.Y. 2023), sanctions for submission of fabricated judicial citations generated by an artificial intelligence tool.

Uniform Guidelines on Employee Selection Procedures, 29 C.F.R. Part 1607.