ANALYST-LED INVESTIGATIVE INTELLIGENCE AND DUE DILIGENCE

RISK INTELLIGENCE

INTELLIGENCE • EVIDENCE • IMPACT

INVESTIGATIVE INTELLIGENCE ASSESSMENT

TCRI-RA-002

Automated Hiring Systems

Applicant Tracking Systems, AI Screening Tools, and Applicant-Flow Outcomes

Open-source investigation into whether automated systems, rather than individualized human review, determine which job applicants receive substantive consideration. Assessment of documented ATS capabilities, AI screening tools, employer configuration, and public hiring-funnel evidence across five enterprise platforms and nine named employers.

SUBJECT Automated hiring systems: ATS platforms and AI screening tools	**SERVICE CATEGORY** Risk Analysis - RA
ANALYTICAL POSTURE Active and Partially Resolved	**INFORMATION CUTOFF** 27 July 2026
PREPARED BY The Collective Risk Intelligence LLC	**HANDLING** Public Distribution

Prepared from vendor documentation, government records and guidance, employer disclosures and bias audits, federal court filings, and peer-reviewed research. Not commissioned by any entity named in this report. No allegation of misconduct is made or implied.

Executive Takeaway

The evidence establishes that every major hiring platform examined provides employer-configurable functions capable of rejecting, disqualifying, qualifying, ranking, or deprioritizing applicants. Several documented functions can operate before substantive human review, depending on employer configuration and workflow. Multiple named employers have deployed AI-enabled scoring or grading tools that classify applicants into priority tiers before recruiter attention. However, the evidence also establishes that employer configuration, not vendor default, determines whether and how these capabilities operate for any given requisition, and that the degree of actual human review before practical exclusion remains the central unresolved question. Public funnel data, where available, confirms steep contraction between application and interview, but the specific stage at which automation versus human judgment drives that contraction cannot be determined from public sources alone. The most consequential gap is the absence of requisition-level data connecting AI scores, recruiter access events, and applicant outcomes.

INVESTIGATION AT A GLANCE

62 SOURCE RECORDS LOGGED	**6** INTELLIGENCE GAPS	**9** OPEN COLLECTION REQUIREMENTS
5 ENTERPRISE ATS PLATFORMS	**9** NAMED EMPLOYERS	**27 JULY 2026** INFORMATION CUTOFF

Dimension	Detail
Investigation type	Open-source intelligence investigation
Subject	Automated hiring systems: ATS platforms, AI screening tools, employer configuration, and applicant-flow outcomes
Scope	U.S. enterprise hiring technology across five major ATS platforms, multiple AI tools, nine named employers, federal hiring systems, and one active federal class-action case
Time period examined	Public records published or updated through July 27, 2026; underlying data spans approximately 2020–2026
Relevant jurisdictions	United States (national); New York City (Local Law 144); Northern District of California (Mobley v. Workday); federal government hiring (OPM/EEOC)
Primary entities	Workday, Greenhouse, Oracle/Taleo, SAP SuccessFactors, iCIMS, HiredScore/Workday, Eightfold AI, Pfizer, Morgan Stanley, TIAA, Citizens Bank, Insmed, McGraw Hill, SharkNinja, DeepHealth, SP+/Bags, U.S. OPM, U.S. Census Bureau, U.S. EPA, iTutorGroup
Primary systems and technologies	Workday Recruiting, HiredScore Spotlight, Candidate Skills Match, Greenhouse Auto-Reject and Talent Matching, Oracle Taleo prescreening, SAP SuccessFactors automated screening, iCIMS Candidate Ranking, USA Staffing, USAJOBS
Source categories	Vendor documentation (Tier 1), government records and guidance (Tier 1), employer disclosures and bias audits (Tier 1–2), court filings and orders (Tier 1–2), ATS aggregate reports (Tier 2), peer-reviewed research (Tier 2), employer career pages (Tier 1)
Collection period	July 26–27, 2026 (primary collection); sources themselves span 2019–2026

Table of Contents

SECTION 01

Purpose

This report presents the results of an open-source intelligence investigation into the role of applicant tracking systems, automated screening tools, and artificial intelligence in modern hiring processes. It distinguishes documented platform capability from employer configuration, separates allegation from adjudicated finding, and states plainly what the public record does not establish. Analyst Comments carry interpretation and are kept visibly separate from the evidentiary record. The investigation was not initiated to prove that any platform, employer, or technology causes discrimination, and it does not assess whether any applicant experienced unlawful discrimination.

SECTION 02

Scope and Limitations

Scope

This investigation examined the publicly documented capabilities, configurations, and outcomes of automated hiring systems across the following dimensions:

- Entities: Major enterprise ATS platforms (Workday, Greenhouse, Oracle Taleo, SAP SuccessFactors, iCIMS), AI screening tools (HiredScore, Eightfold AI, Candidate Skills Match, iCIMS Candidate Ranking, Greenhouse Talent Matching), named employers deploying these tools (Pfizer, Morgan Stanley, TIAA, Citizens Bank, Insmed, McGraw Hill, SharkNinja, DeepHealth, SP+/Bags, and employers named in Mobley v. Workday), LinkedIn platform documentation used only for applicant-count definitions and Recruiter screening functions, and federal hiring systems operated through USA Staffing and USAJOBS [1].
- Timeframe: Public sources published or updated through July 27, 2026. Employer disclosures, bias audits, and litigation filings span periods from approximately 2020 through mid-2026. Federal applicant-flow data spans fiscal years 2023 through 2025. ATS aggregate benchmarks cover January 2021 through March 2026 [1].
- Jurisdictions: United States, with particular attention to New York City (Local Law 144), federal government (EEOC MD-715, OPM staffing systems), and Northern District of California (Mobley v. Workday).
- Source categories: Official vendor documentation, government regulatory guidance, government system documentation, employer-published bias audits and AEDT notices, federal court filings and orders, agency applicant-flow disclosures, ATS aggregate benchmark reports, peer-reviewed research, and employer career pages and job postings [1].

Out-of-Scope Matters

The following related issues were intentionally excluded:

- Individual applicant qualification assessments or case-by-case employment outcomes.
- Employer-specific ATS configuration settings, internal recruiter workflows, or proprietary threshold data not publicly disclosed. AI systems used outside hiring, such as employee performance monitoring or internal mobility tools, except where they share platform architecture with recruiting tools.
- International hiring regulations or non-U.S. employer practices.
- U.S. state and local AI-hiring regimes other than New York City Local Law 144, including the Illinois Artificial Intelligence Video Interview Act, the Colorado Artificial Intelligence Act, and California Fair Employment and Housing Act regulations governing automated-decision systems, together with changes in federal enforcement posture toward AI in hiring. These regimes are live and relevant to the subject. Local Law 144 was selected because it is the only regime that compels public bias-audit disclosure, which is what makes employer-level evidence collectable; the others impose notice, assessment, or consent duties that generate no comparable public record.
- Speculation about unreported employer behavior or undisclosed system features.
- Any assessment of whether specific applicants experienced unlawful discrimination.

UNAVAILABLE PRIVATE RECORDS: The most consequential limitation is that the internal configuration of employer ATS settings, scoring thresholds, queue order, recruiter review depth, and actual disposition logic are almost never publicly available. Vendor documentation establishes what a system can do; it does not establish what any employer chose to do for a given requisition [1].

INCOMPLETE PUBLIC REPORTING: Federal applicant-flow data requires reconciliation across multiple source systems before analytical use, as documented by EPA's FY 2025 MD-715 report [22]. Interview-stage counts are often missing from published federal hiring data, even though EEOC MD-715 instructions explicitly require their reporting

[27]. Private-sector employers are under no general obligation to publish applicant-flow data, making requisition-level funnel evidence extremely scarce outside the federal sector.

INCONSISTENT TERMINOLOGY: Platforms, employers, and public sources use different definitions for "applicant," "application," "rejection," "screening," and "review." LinkedIn's applicant count for external-site applications reflects clicks to apply, not completed applications [51]. Separate LinkedIn documentation defines how applications are counted and reported to job posters, how applicants flow into Recruiter, and how free and promoted postings differ [48, 47, 50, 46]. ATS aggregate benchmarks measure applications per hire rather than applications per posting [30]. These definitional differences prevent direct cross-source comparisons without careful reconciliation [1].

LITIGATION EVIDENCE CONSTRAINTS: Court filings in Mobley v. Workday contain unverified allegations, procedural rulings, and discovery disputes, not adjudicated findings of fact. Preliminary collective certification is not a merits determination [8]. Alleged audit data covering 724,352 applicants has not been independently verified or produced publicly [5]. The court denied plaintiffs' motion to compel Workday's bias-testing data on privilege grounds and found plaintiffs had not demonstrated Workday's legal control over customer applicant data [6].

BIAS AUDIT INTERPRETIVE LIMITS: Published NYC Local Law 144 bias audits vary materially in data type, pooling methodology, employer-specific coverage, and missing demographic data. Peer-reviewed research found that 75% used historical data, 16% used test data, and 9% did not identify the data type, and that pooling and silent duplicates often made employer-specific interpretation difficult [37].

UNRESOLVED ATTRIBUTION: For most employer deployments, the publicly available evidence cannot distinguish whether an observed outcome (rejection, non-advancement, rapid disposition) was caused by the ATS platform's common architecture, the employer's specific configuration, the applicant's qualifications relative to the posting, recruiter discretion, market competition, or some combination of these factors.

SECTION 03

Objectives and Intelligence Requirements

The investigation was not initiated to prove that any particular platform, employer, or technology causes discrimination, nor to validate any applicant's individual claim. Rather, it was conducted to answer a set of neutral intelligence questions using verifiable public sources, and to clearly distinguish what the evidence supports from what remains unknown, alleged, or unresolved.

Objective

The investigative objective is to produce a defensible, evidence-grounded assessment of how automated hiring systems operate across major enterprise ATS platforms and AI-enabled screening tools, and to identify the publicly available evidence concerning applicant volume, funnel contraction, automated disposition, human review, and the distribution of decision-making authority between software and people. The intended analytical outcome is a structured intelligence product that decision-makers, whether regulators, applicants, employers, researchers, or legal practitioners, can use to understand what is publicly documented, what remains opaque, and where the most consequential intelligence gaps exist.

Primary Intelligence Question

To what extent do publicly documented ATS capabilities, AI screening tools, employer configurations, and hiring-funnel data establish whether automated systems, rather than individualized human review, determine whether job applicants receive substantive consideration?

Supporting Intelligence Questions

The following secondary questions are required to answer the primary intelligence question. Each map to one or more collection requirements established during the investigation [1]:

- What automatic reject or disqualify functions does each major ATS publicly document? (CR-ATS-001)
- Who can create, activate, modify, or inherit the screening rules that these systems execute? (CR-ATS-002)
- Can applications be ranked, filtered, scored, hidden, or prioritized without formal rejection, and how does this affect practical consideration? (CR-ATS-003)
- What evidence can establish whether a human reviewed an application before disposition? (CR-ATS-004)
- What does each job platform count and display as an "applicant," and how do these definitions affect reported volumes? (CR-VOL-001)
- What public data exists on applications, recruiter review, interviews, offers, and hires at the requisition or role level? (CR-VOL-002)
- What AI or external screening tools can influence candidate status, ranking, or visibility inside the ATS workflow? (CR-AI-001)
- What evidence shows recruiters or hiring teams use AI tools after ATS screening or outside the ATS workflow? (CR-AI-002)
- What legal, regulatory, procurement, or court records disclose actual hiring-system configuration or outcomes? (CR-LEGAL-001)

These questions were framed as neutral collection requirements, not propositions to confirm. For each question, the collection log records supporting, contradicting, and limiting evidence with equal specificity [1].

SECTION 04

Investigative Summary

The primary intelligence question asked to what extent public documentation establishes that automated systems, rather than individualized human review, determine whether applicants receive meaningful consideration.

The investigation reviewed and validated 62 unique public source records, each numbered in Section 17, including official vendor documentation from Workday, Greenhouse, Oracle Taleo, SAP SuccessFactors, and iCIMS; employer-published bias audits and AEDT disclosures from Pfizer, Morgan Stanley, TIAA, Citizens Bank, and others; federal court filings and orders in Mobley v. Workday; EEOC enforcement records; federal applicant-flow data from OPM, Census, and EPA; ATS aggregate benchmarks from Ashby and CareerPlug; and peer-reviewed research on bias audit quality [1].

Most significant findings:

- Every major hiring platform examined publicly documents at least one employer-configurable function capable of rejecting, disqualifying, qualifying, ranking, or prioritizing applicants [62, 58, 55, 53, 52, 9, 49].
- AI-enabled prioritization tools, most notably Workday HiredScore Spotlight, assign each application within the covered population a match grade (A through D) that structures which candidates recruiters see first. Pfizer's bias audit data shows approximately 56.5% of audited applications from applicants with known gender fell outside the A/B priority grades [14]. No public evidence establishes whether C/D-graded applicants receive equivalent review.
- Named employers publicly confirm that AI tools score, rank, or recommend candidates before human review, but the specificity and enforceability of their human-review commitments vary materially, from SharkNinja's requirement that a trained recruiter review the full application before any shortlisting decision [13] to SP+/Bags' conditional disclosure that it "may utilize" an AEDT with no stated review standard [11].
- Public funnel data confirms steep contraction: Ashby reports 291 applications per hire in early 2026 across its customer base, and CareerPlug reports approximately 3% of applicants receive interview invitations [30, 28]. Federal data from Census shows measurable contraction from 225 GS-15 applications to 6 selections, with no interview-stage data published [26].
- The Mobley v. Workday litigation provides the first cross-employer outcome-pattern allegation tied to a single ATS platform, with the court granting preliminary collective certification based on a sufficiently alleged unified Workday AI scoring policy. However, this remains a procedural finding, not a merits determination [8].

OVERALL ASSESSMENT: The evidence strongly supports the conclusion that automated systems play a material role in determining which applicants receive substantive human consideration, and that employer configuration rather than vendor default is the primary operational determinant of automated disposition. Whether the resulting outcomes are discriminatory, job-related, or lawful remains outside the evidentiary reach of public sources and is not resolved by this investigation.

CONFIDENCE LEVEL: High confidence that the documented capabilities exist and are deployed; moderate confidence that AI prioritization materially affects practical consideration across employer implementations; low confidence in any specific platform-wide rejection rate, because the data needed to calculate one has not been publicly produced or independently verified.

PRINCIPAL INTELLIGENCE GAPS: Requisition-level data connecting AI scores or grades with recruiter access events, interview decisions, and final outcomes. Employer-specific threshold, queue-order, and review-depth documentation. The underlying data from the alleged Workday 724,352-applicant audit. Interview-stage counts from both public and private employers.

SECTION 05

Analytical Findings

1 Major Hiring Platforms Publicly Document Automated Rejection, Disqualification, Qualification, or Prioritization Without Human Review at the Point of System Action

ASSESSMENT: The major hiring platforms examined provide officially documented automated functions that can reject, disqualify, qualify, rank, or prioritize applicants based on employer-configured criteria without requiring individualized human review at the time the system executes the action. The specific function differs by platform and should not be treated as a uniform rejection mechanism.

SUPPORTING EVIDENCE: Greenhouse documents auto-reject rules triggered by custom application question responses, with configurable rejection reasons and automated applicant emails [62, 61]. Oracle Taleo documents mandatory disqualification questions that determine whether candidates move forward or are automatically disqualified [58]. SAP SuccessFactors documents questions that can be marked as disqualifiers, automatically disqualifying applicants who do not provide the correct answer, with the system placing the application into an "Automatically Disqualified" status [55, 53]. Workday markets conversational AI that asks required knockout questions and qualifies candidates using employer-defined criteria, moving candidates forward automatically [52]. iCIMS documents Candidate Ranking, which provides recruiters a ranked list of applicants against position requirements [9]. LinkedIn Recruiter permits auto-rejection of applicants who do not meet required screening qualifications [49].

Additionally, Greenhouse documents that after a candidate receives a security-concern rejection, all future applications to the organization are automatically rejected without additional review [59]. SAP SuccessFactors documents configurable email delays for disqualification notifications that apply equally to manual, automatic, and API-based disqualification actions, meaning the timing of a rejection email does not reliably distinguish whether the decision was human or automated [54].

CONTRADICTORY OR QUALIFYING EVIDENCE: Greenhouse stated in 2023 that its rules-based automation lets employers define qualifications while "people determine whether candidates meet them," and characterized its product as outside NYC's AEDT scope [10]. This statement predates Greenhouse's 2026 Talent Matching deployment and does not describe every current feature [1]. Capability does not establish that any employer enabled or used a particular feature for a particular requisition.

EVIDENTIARY LIMITATIONS: The documentation establishes what each platform can do. It does not establish the frequency of employer activation, the content of the rules applied, or whether any specific applicant was disposed of by automation rather than human decision.

CONFIDENCE LEVEL: High.

> ANALYST COMMENT
>
> This finding is foundational. The automated functions described in this finding are not allegations or inferences. Each platform officially documents at least one mechanism for automated rejection, disqualification, qualification, ranking, or prioritization. The critical analytical step is separating this established capability from the separate question of whether it was activated, correctly configured, and applied to a particular applicant at a particular employer. The evidence for that second question is almost entirely unavailable from public sources, which is itself a significant intelligence gap.

2 AI-Enabled Prioritization Tools Classify Applicants Into Tiers That Structure Recruiter Attention, but No Public Evidence Establishes That Lower-Priority Applicants Receive Equivalent Review

ASSESSMENT: Multiple named employers deploy AI scoring or grading tools that assign each application within the covered or audited population a comparative match designation before recruiters begin review. The most publicly documented is Workday HiredScore Spotlight, which assigns A through D grades against employer-defined requisition requirements. The evidence supports the conclusion that these tools materially influence which candidates receive recruiter attention, but the evidence does not establish a specific rejection rate or prove that lower-graded candidates are systematically excluded.

SUPPORTING EVIDENCE: Pfizer's published bias audit reports that HiredScore assigns each application within the audited population an A–D grade, with approximately 43.5% of applications from applicants with known gender receiving A or B grades, meaning approximately 56.5% fell into C or D grades [14]. TIAA's pre-deployment audit used a binary priority classification of A/B versus C/D grades [39]. Morgan Stanley's audit of the Eightfold model analyzed 956,931 applications with match scores from 0.5 to 5, considered by recruiters when deciding whom to advance [40]. Citizens Bank's audit shows its AEDT classifies assessment scores into bottom, middle, and top tiers using percentile-ranked scores grouped into thirds [41]. Workday's product page states HiredScore provides "accelerated candidate prioritization" and "guidance to high-value recruiter tasks" integrated with Workday Recruiting [15].

CONTRADICTORY OR QUALIFYING EVIDENCE: Pfizer states that "recruiters and hiring managers always have complete discretion over which applicants they select and what information they consider" and that "HiredScore does not automatically make or execute any decisions" [14]. TIAA expressly reserved whether the Spotlight tool legally qualifies as an AEDT [39]. Workday publishes its own Spotlight analysis for transparency but states it is not intended to satisfy customer-specific obligations [2].

EVIDENTIARY LIMITATIONS: No public source establishes the actual review rate, access event, or progression outcome by grade or tier. A grade distribution is not a rejection rate. The audits report what the tool outputs; they do not report what happened to applicants after the output was generated. The critical gap is post-grading recruiter behavior data.

CONFIDENCE LEVEL: High that the grading and prioritization occur at documented scale; moderate that they materially affect practical consideration; low that any specific non-review rate can be calculated from public data.

> ANALYST COMMENT
>
> The design purpose of these tools is to direct recruiter attention to the highest-match candidates. When an employer processes hundreds of thousands of applications (as documented in the Morgan Stanley and Pfizer audits), it is analytically implausible that every C/D-graded or low-scored application receives the same depth of review as an A-graded application. But plausibility is not proof. The evidence needed to convert this analytical inference into a documented finding (recruiter access logs, per-grade interview rates, and outcome data) is not publicly available.

3 Employer Configuration, Not Vendor Default, Determines How Automated Screening Operates for Each Requisition

ASSESSMENT: Across all platforms and tools examined, the evidence consistently demonstrates that the specific criteria, thresholds, questions, and workflow rules that determine automated disposition are configured by the employer or its authorized users, not imposed by the ATS vendor as a fixed default.

SUPPORTING EVIDENCE: Greenhouse's auto-reject documentation requires the employer to create a custom job post question and select the response that triggers rejection [62]. Oracle Taleo assigns disqualification-question creation responsibility to authorized customer content managers [56]. SAP SuccessFactors allows questions to be marked as disqualifiers at the requisition level, with the employer configuring correct answers [55]. Workday's discovery responses in Mobley describe Candidate Skills Match as a feature that customers "can choose to turn on or off" [8]. The iTutorGroup EEOC settlement confirmed that the employer programmed its application software to automatically reject applicants based on age thresholds, demonstrating that employer-defined rules can independently create a high automatic rejection rate without evidence that the ATS vendor supplied the discriminatory criterion [3, 4].

CONTRADICTORY OR QUALIFYING EVIDENCE: Plaintiffs in Mobley allege that Workday's AI recommendation system operates through "standardized mechanisms" across employers and that "the core discriminatory mechanisms are common and centrally designed by Workday" [5]. The court found this sufficient for preliminary collective certification [8]. However, the court also noted that "how Workday's AI recommendation system works, and the extent to which any disparate impact is generally driven by employer preferences rather than Workday's own algorithms, are issues susceptible to common proof" at a later stage, not issues resolved at the certification stage [8].

EVIDENTIARY LIMITATIONS: The evidence establishes the governance model (employer configures, platform executes) but does not establish how much latitude each employer exercises versus how much is constrained by platform defaults, templates, or inherited settings. Some platforms offer centralized question libraries that may propagate outdated criteria across requisitions [56].

CONFIDENCE LEVEL: High.

> ANALYST COMMENT
>
> The iTutorGroup case is the strongest control in this investigation. The EEOC documented more than 200 qualified applicants automatically rejected because the employer programmed age-based knockout rules into its application software [3, 4]. No vendor default caused this outcome. This does not mean vendors bear no responsibility for platform design choices, but it demonstrates that the proximate cause of automated rejection in documented cases has been employer configuration. The Mobley litigation tests whether platform-level architecture itself produces discriminatory patterns independent of any single employer's settings, a question that remains unadjudicated.

4 Public Hiring-Funnel Data Confirms Steep Contraction Between Application and Interview, but the Stage at Which Automation Versus Human Judgment Drives That Contraction Cannot Be Determined from Public Sources

ASSESSMENT: Multiple independent data sources confirm that applicant pools contract sharply between submission and interview, with selection rates typically below 5% and interview invitation rates as low as 3%. However, no public source establishes the precise stage at which automated screening, AI prioritization, or human recruiter judgment is the primary cause of contraction for any specific employer or platform.

SUPPORTING EVIDENCE: Ashby reports that applications per hire averaged 291 in early 2026 across its 109-million-application dataset, and that candidates are roughly 50% less likely to receive an interview today than in 2021 [30, 29]. CareerPlug reports an applicant-to-interview ratio of approximately 3% across its 2024 benchmark dataset [28]. The U.S. Census Bureau published actual stage counts: for GS-15 PWD new hires, 225 applicants, 209 qualified, 165 referred, and 6 selected (a 2.7% selection rate); for GS-13 PWD new hires, 23 applicants, 9 qualified, 6 referred, and 1 selected [26]. Workday reported that job applications grew four times faster than job openings in the first half of 2024 [44].

CONTRADICTORY OR QUALIFYING EVIDENCE: These datasets measure different things. Ashby's metric is applications per hire, not applications per posting. CareerPlug primarily serves small and midsize employers. Census data is grouped by disability status and grade category and may combine multiple requisitions. LinkedIn's external-site applicant count reflects clicks to apply, not completed applications [51]. No interview-stage count was published in the Census data. The two Census cohorts also cut against a pure volume explanation: the smaller GS-13 pool shows a higher selection rate (4.3%) than the larger GS-15 pool (2.7%) [26].

EVIDENTIARY LIMITATIONS: The public evidence establishes that steep funnel contraction is widespread, but it cannot attribute the contraction to automation, human screening, applicant qualification, market competition, or some combination. The federal hiring system has the most granular publicly documented funnel structure (applied → qualified → referred → interviewed → selected), but even there, interview data is frequently missing from published agency reports [27, 22].

CONFIDENCE LEVEL: High that steep funnel contraction occurs; low that public data can attribute the contraction to any single cause.

> ANALYST COMMENT

> The funnel data is significant because it establishes the scale of the screening problem that ATS automation and AI tools are designed to address. When an employer receives 291 applications per hire, some mechanism must reduce the pool before meaningful human review can occur. The question is not whether filtering happens, since it must, but whether the filtering mechanism is lawful, job-related, and transparent. Public data cannot answer that question at the employer or requisition level.

5 The Mobley v. Workday Litigation Provides the First Cross-Employer Outcome-Pattern Allegation Tied to a Single ATS Platform, but Remains at an Early Procedural Stage with No Adjudicated Factual Findings

ASSESSMENT: The Mobley case is the most significant active litigation concerning platform-level automated hiring discrimination. It provides unique evidence about Workday's product architecture (Candidate Skills Match categories, HiredScore integration, and assessment tools) and includes allegations of a large cross-customer audit. However, all substantive allegations remain contested and unadjudicated.

SUPPORTING EVIDENCE: The named plaintiff alleges more than 100 applications through Workday employers and rejection for every application. Four opt-in plaintiffs allege hundreds of applications and rejection nearly every time without interviews [8]. The court granted preliminary collective certification after finding a sufficiently alleged unified policy involving Workday's AI recommendation system to "score, sort, rank, or screen applicants" [8]. Workday's discovery responses described Candidate Skills Match as reporting results as "strong," "good," "fair," "low," "pending," and "unable to score" [8]. The third amended complaint alleges that a Workday-commissioned external audit analyzed 724,352 applicants across ten large customers and found statistically significant disparate impact by race and gender [5]. The EEOC filed an amicus brief arguing that software may determine whether applicants are "referred to and considered by employers" [7]. The discovery order confirmed the existence of bias-testing data, anonymized application and resume data, and customer applicant data stored on Workday servers [6].

CONTRADICTORY OR QUALIFYING EVIDENCE: The court expressly stated that preliminary certification "is not a merits finding and does not establish disparate impact or causation" [8]. Workday asserted in discovery that it "does not recommend, screen out, or otherwise assess or predict applicants' likelihood of success in a role" [8]. The court denied plaintiffs' motion to compel production of Workday's bias-testing data, finding it protected by attorney-client privilege [6]. The court also found that plaintiffs had not shown Workday controls customer applicant data for Rule 34 production purposes [6]. The alleged 724,352-applicant audit has not been publicly produced, independently verified, or entered as evidence [5].

EVIDENTIARY LIMITATIONS: The complaint contains allegations, not proven facts. The underlying audit data is not public. The named employer cohort (Uline, Cigna, GE HealthCare, Cardinal Health, Ciena, Intel, Stryker, Thermo Fisher Scientific) has not been individually profiled for ATS configuration, HiredScore activation, or rejection-cause documentation [5]. Employer criteria, applicant qualifications, market competition, and individual recruiter decisions remain alternative explanations for each reported rejection.

CONFIDENCE LEVEL: High for procedural and product-function facts documented in court orders; moderate for the existence of a cross-employer pattern allegation; low for any validated conclusion about platform-level causation or discrimination.

> ANALYST COMMENT
>
> Mobley is analytically significant because it is the only case that tests whether a common ATS platform architecture, as opposed to an individual employer's configuration, can itself produce legally cognizable discrimination. The EEOC's amicus participation signals federal enforcement interest in the platform-as-gatekeeper theory. But the distinction between the procedural posture (notice may issue to a potentially vast collective) and the merits (has Workday caused discriminatory outcomes?) is enormous. The investigation treats Mobley as a high-value collection lead for understanding Workday's product architecture, not as proof of any conclusion about discrimination.

SECTION 06

Definitions and Analytical Controls

The following definitions and controls govern how evidence, terminology, and source claims are used throughout this report. They are designed to prevent confirmation bias, unsupported attribution, and inconsistent interpretation of terms that carry different meanings across platforms, employers, and legal contexts.

Term	Working Definition	What Must Not Be Inferred
Automated Rejection	A system action that changes an applicant to a rejected or disqualified status without a user manually selecting that applicant at the time of disposition [1].	Do not infer that the ATS vendor independently selected the rejection criterion. Identify the rule creator, trigger, system action, and available override.
Employer-Configured Automation	A rule, question, threshold, workflow, or integration selected or configured by an employer that the system executes [1].	Do not attribute the criterion to the ATS vendor unless vendor documentation shows it is a fixed default. Separate platform capability from customer configuration.
Human Review	Documented evidence that a person accessed and substantively evaluated application content [1].	A status change, email delay, or record-open event alone does not prove substantive review. Seek audit fields, workflow documentation, or public records defining review actions.
Applicant Count	Completed applications for a specific requisition, unless the source explicitly defines another metric [1].	Do not treat clicks, redirects, starts, profiles, or expressions of interest as completed applications. Record each platform's counting method.
Prioritization vs. Rejection	AI-generated grades, scores, or rankings that order or classify applicants are treated as prioritization, not rejection, unless the source establishes that lower-priority applicants are hidden, excluded, or automatically rejected [1].	Do not equate a C/D grade or "low" score with automatic rejection unless the employer's workflow documentation establishes that outcome.
Allegation vs. Finding	Court pleadings, complaints, and party assertions are allegations. Adjudicated rulings on the merits, government enforcement settlements, and independently verified data are findings [1].	Preliminary collective certification is not a merits finding. An alleged audit is not a validated dataset.
Selection Rate vs. Rejection Rate	A scoring rate or grade distribution is not a selection or rejection rate unless the source ties the score or grade to an actual advancement or rejection decision [1].	Do not convert Pfizer's A/B grade distribution into a "rejection rate" or treat Morgan Stanley's scoring-rate data as an interview or advancement rate.

SECTION 07

Analytical Framework

The investigation uses a layered decision-delegation framework to separate the distinct stages at which technology, configuration, and human judgment affect applicant outcomes. This framework is adapted to the specific evidence collected and is not a generic hiring-process model.

PLATFORM CAPABILITY LAYER: What the ATS or AI tool can do, as documented by the vendor. This includes auto-reject rules, prescreening questions, AI scoring, candidate grading, and automated status changes. Evidence at this layer comes from vendor documentation, product pages, and official demonstrations [1].

EMPLOYER CONFIGURATION LAYER: What the employer chose to activate, configure, and apply to specific requisitions. This includes the content of knockout questions, the enabling of AI features like CSM or HiredScore, the definition of requisition requirements, and the setting of any thresholds or referral limits. Evidence at this layer comes from employer AEDT notices, bias audits, job postings, and court-produced configuration descriptions [1].

AUTOMATED EXECUTION LAYER: The system actions that occur without a human manually selecting the affected applicant: auto-reject, auto-disqualify, auto-advance, or automatic status change. Evidence comes from vendor test scripts, workflow documentation, and employer disclosures [53, 62].

AI PRIORITIZATION LAYER: The assignment of scores, grades, rankings, or classifications that determine the order or grouping in which applications appear to recruiters. This layer operates between automated execution and human review. Evidence comes from bias audits, vendor product descriptions, and court-cited discovery responses [14, 40, 8].

RECRUITER VISIBILITY AND ACCESS LAYER: Whether and how recruiters see, open, and evaluate applications. This is the most significant intelligence gap. Almost no public evidence documents recruiter access events, review depth, or the practical effect of AI prioritization on which applications are actually opened. Evidence at this layer would come from audit logs, access-event records, or employer-specific workflow documentation, none of which are publicly available for any private employer examined.

HUMAN DECISION LAYER: The point at which a person makes a substantive judgment about an applicant's qualifications and decides to advance, interview, or reject. Evidence comes from employer process descriptions, such as Citizens Bank's statement that "recruiters review your application and assessment" [17] or SharkNinja's statement that "a trained recruiter reviews the tool output and the full application before any screening or shortlisting decision" [13].

OBSERVABLE OUTCOME: The end-state visible to the applicant or to public reporting: rejection email, status change, interview invitation, offer, or hire. Evidence comes from applicant-facing communications, funnel data, and court-cited rejection notices.

The critical analytical question is at which layer the most consequential filtering occurs for any given employer. Public evidence is strongest at the Platform Capability and Automated Execution layers and weakest at the Recruiter Visibility and Human Decision layers.

SECTION 08

Factual Background

The factual context of this investigation spans three interconnected domains: ATS platform capabilities, AI-enabled hiring tools, and the regulatory and legal environment in which they operate.

ATS PLATFORM EVOLUTION: The five platforms examined, Workday Recruiting, Greenhouse, Oracle Taleo, SAP SuccessFactors, and iCIMS, each provide employer-configurable automated screening capabilities. These capabilities have existed in some form for years (Taleo's prescreening documentation references question libraries managed by content managers [56]; Greenhouse's rules-based automation was described in 2023 [10]), but have been augmented since approximately 2023–2024 by integrated AI tools. Workday acquired HiredScore in April 2024 and integrated its AI-powered candidate prioritization, talent discovery, and recruiter guidance features into Workday Recruiting [15, 5]. Greenhouse Talent Matching, which ranks applications against role requirements using CV content, is documented in employer AEDT disclosures rather than in any Greenhouse public product documentation located during collection [13, 1]. iCIMS deployed Candidate Ranking, which provides recruiters a ranked list of applicants against position requirements [9].

SCALE OF THE HIRING INFRASTRUCTURE: Workday services more than 10,000 organizations, including over half of Fortune 500 companies [5]. The Mobley complaint alleges that in May 2023 alone, Workday processed approximately 2.2 million job requisition transactions [5]. Workday reported processing 19 million job requisitions and 173 million applications in the first half of 2024 [44]. Ashby's dataset covers 109 million applications and 247,000 jobs from January 2021 through March 2026 [29].

NYC LOCAL LAW 144: Enacted December 2021 and enforced from July 5, 2023, Local Law 144 requires employers using automated employment decision tools in New York City to publish annual bias audit summaries [42, 43]. The law defines AEDTs as tools that "substantially assist or replace discretionary decision making" in hiring [42]. Required disclosures include selection or scoring rates, impact ratios, data source and explanation, and unknown-category counts [43]. The law does not require any specific action based on audit results, does not mandate that impact ratios meet any threshold, and does not require publication of individual decisions or thresholds [43]. Peer-reviewed analysis of 44 reports containing 116 audits found significant limitations including missing demographic data, data pooling, test data use, and silent duplicates [37].

FEDERAL HIRING FUNNEL STRUCTURE: The federal government maintains the most granular publicly documented hiring funnel through OPM's USA Staffing system and EEOC's MD-715 reporting requirements. OPM's data model includes application-level indicators for qualified, not referred–not qualified, not referred–qualified, referred, selected, sent tentative offer, sent official offer, hired, and entered on duty [25]. EEOC instructs agencies to report the number of people who applied, were qualified, referred, interviewed, and selected [27]. The 2025 Merit Hiring Plan requires agencies to identify the methodology for determining how many applicants will be referred and permits cut-off scores, set numbers, and percentage-based referral methods [32, 31]. OPM has separately maintained an applicant-flow data workgroup and governs agency access to applicant-flow data [23, 35], and USAJOBS publishes indicative timelines for each stage of the federal hiring process [24].

ENFORCEMENT PRECEDENT: In 2023, iTutorGroup settled an EEOC enforcement action for $365,000 after the agency alleged the employer programmed its online application software to automatically reject female applicants age 55 or older and male applicants age 60 or older, rejecting more than 200 qualified applicants [3, 4].

SECTION 09

Evidence

Documentary Evidence

Official vendor documentation across the five ATS platforms establishes automated rejection, disqualification, qualification, prescreening, or ranking capabilities at the requisition level. Greenhouse documents three application rules (auto-reject, auto-advance, auto-tag) configurable per job post [61, 62, 60]. Oracle Taleo documents mandatory disqualification questions in a library managed by content managers [58, 56], and publishes guidelines for scoring prescreening questions [57]. SAP SuccessFactors documents requisition-level disqualifier questions with automatic disqualification and a configurable email delay that applies identically to manual, automatic, and API-based actions [55, 53, 54, 36]. Workday documents knockout questions and automated candidate qualification through conversational AI [52]. iCIMS documents Candidate Ranking providing a ranked applicant list against position requirements [9].

Employer-published AEDT notices and bias audits provide the most specific evidence of deployed AI tool use. Pfizer's audit covers all applicants to NYC-hiring requisitions from August 2022 through February 2025 using HiredScore A–D grading [14]. Morgan Stanley's audit covers 956,931 applications scored by Eightfold's matching model from June 2024 through May 2025 [40]. TIAA's audit covers Workday-provided data with A/B versus C/D binary priority classification [39]. SharkNinja's notice describes Greenhouse Talent Matching with a stated human-review requirement before any screening decision [13]. DeepHealth identifies iCIMS as its AEDT for certain NYC postings [12]. Insmed discloses AI-based resume scoring in a Workday-centered hiring environment [21, 20].

These documents establish that AI tools are actively deployed at documented scale, but they do not establish disposition thresholds, review depth, or actual applicant outcomes by score or grade [1].

REGULATORY AND LEGAL EVIDENCE, MOBLEY V. WORKDAY: The Workday court record provides the most detailed publicly available description of Workday's product architecture. The court's May 2025 order describes Candidate Skills Match reporting results as "strong," "good," "fair," "low," "pending," or "unable to score," and notes that employers can enable, disable, use, or ignore the feature [8]. The court granted preliminary collective certification based on an alleged unified policy of AI-enabled scoring, sorting, ranking, or screening [8]. The May 2026 discovery order confirmed the existence of bias-testing data and anonymized application data on Workday servers, denied production of privileged bias-testing data, and found plaintiffs had not demonstrated Workday's control over customer applicant data [6]. The July 2026 order adopted strikes to the third amended complaint and directed filing of a fourth amended complaint [5].

The EEOC's amicus brief argued that screening software may control access to employer consideration and that Workday could be liable as an employment agency, indirect employer, or agent of employers [7]. The iTutorGroup settlement established that employer-programmed automatic rejection based on discriminatory criteria constitutes unlawful discrimination enforceable by the EEOC [3, 4].

NYC DCWP guidance defines the AEDT scope, bias audit requirements, notice obligations, and data types, but expressly states that the law does not require any specific action based on audit results [42, 43].

Public funnel data exists but is fragmented and methodologically inconsistent across sources:

Source	Period	Population	Applications	Qualified	Referred	Interviewed	Selected or Hired	Reported Metric
U.S. Census Bureau [26]	FY 2023	GS-15 PWD new hires	225	209	165	—	6	2.7%
U.S. Census Bureau [26]	FY 2023	GS-13 PWD new hires	23	9	6	—	1	4.3%
CareerPlug [28]	2024	Cross-industry benchmark	~100	—	—	~3	—	~3% interview rate
Ashby [30]	Early 2026	Cross-customer benchmark	291	—	—	—	1	291 applications per hire

These data points are not directly comparable because they measure different populations, use different denominators, and cover different stages. Interview-stage data is missing from all published datasets in this investigation except the CareerPlug benchmark estimate [1].

Entity and System Evidence

Nine employer technology profiles were developed from publicly available sources [1]. The evidence establishes:

- Workday deployments (Pfizer, TIAA, Insmed): HiredScore or unidentified AEDT scoring within Workday Recruiting, with A–D grading documented for Pfizer and TIAA, and unspecified AI scoring for Insmed [14, 39, 21, 20, 19].
- Greenhouse deployments (SharkNinja): Talent Matching with stated human-review control; planned NYC deployment beginning August 2026.
- iCIMS deployments (DeepHealth, SP+/Bags, McGraw Hill): Platform-level AEDT use identified by DeepHealth; conditional disclosure by SP+/Bags; like-skills recommendation by McGraw Hill [12, 11, 18].
- Standalone AI screening and assessment tools (Morgan Stanley, Citizens Bank): Eightfold matching-model scoring of 956,931 applications at Morgan Stanley, with an independent audit of the same model published by the tool vendor [40, 38]; percentile-tiered assessment scoring at Citizens Bank, with recruiter review of application and assessment stated as the step preceding hiring-manager scheduling [41, 17, 16].
- Cross-employer litigation cohort (Uline, Cigna, GE HealthCare, Cardinal Health, Ciena, Intel, Stryker, Thermo Fisher Scientific): Named in Mobley pleading as Workday customers where plaintiffs applied and were rejected, but individual technology profiles and HiredScore activation are not confirmed from public sources.

Human review language varies materially: SharkNinja commits to full-application review before any screening decision [13]; Pfizer states complete discretion but does not state every application receives review before shortlisting [14]; DeepHealth promises substantive human review before "any final decision" but does not address initial screening [12]; SP+/Bags makes no human-review statement [11].

SECTION 10

Comparative Analysis

Comparison across ATS platform cohorts reveals both structural commonalities and material employer-level variation:

PLATFORM COMMONALITY: All three major platforms examined in depth (Workday, Greenhouse, iCIMS) support software-mediated prioritization, scoring, or ranking of applicants. Workday's HiredScore uses A–D grading [14]; Greenhouse's Talent Matching ranks applications [13]; iCIMS offers Candidate Ranking [9]. This architectural commonality is well-documented.

EMPLOYER VARIATION WITHIN PLATFORMS: Within the Workday cohort alone, Pfizer publishes a complete grade distribution with a statement of recruiter discretion but no stated per-application review standard [14]; TIAA conducted a pre-deployment audit but reserved whether the tool is an AEDT [39]; Insmed discloses AI scoring but does not identify the tool vendor [21]. Within the iCIMS cohort, DeepHealth identifies iCIMS as its AEDT and promises review before final decisions [12], while SP+/Bags uses conditional language ("may utilize") with no review statement [11].

HUMAN REVIEW COMMITMENTS ARE NOT STANDARDIZED: SharkNinja's notice is the most specific, requiring trained-recruiter review of both the tool output and the full application before any screening or shortlisting decision [13]. By contrast, DeepHealth's commitment to "substantive human review and intervention before any final decision" does not address initial screening [12]. Citizens Bank states recruiters review application and assessment before scheduling with hiring managers, but this does not establish that every application reaches that stage [17]. No employer examined publishes evidence of how many applications are actually opened by a human before disposition.

FEDERAL VERSUS PRIVATE-SECTOR FUNNEL TRANSPARENCY: The federal hiring system documents the most granular public funnel structure, with OPM maintaining vacancy-level metrics for applications received, qualified, referred, and selected [34, 25]. EPA reported developing a process to validate applicant-flow data from competing sources, correct errors, and create a dashboard to automate data transfer into MD-715 tables [22]. No comparable private-sector funnel transparency exists in the collected evidence.

COMPARABLE EMPLOYER REJECTION RATES ARE NOT AVAILABLE: Despite the Mobley allegations and the bias audit disclosures, no validated ATS-wide or employer-specific rejection rate attributable to AI scoring has been published. The Mobley complaint alleges a 724,352-applicant audit showing disparate impact, but the underlying data has not been produced or verified [5]. Pfizer's grade distribution is not a rejection rate [14]. The iTutorGroup settlement confirms that employer configuration can cause a high and measurable automatic rejection pattern, but the platform was not identified [3, 4].

SECTION 11

Competing Explanations

The observed evidence, steep funnel contraction, rapid rejections, cross-employer rejection patterns, and AI-mediated prioritization, is consistent with multiple explanations, none of which can be conclusively ruled out from public sources alone.

EXPLANATION 1: Employer-configured automation drives most automated rejections. Supporting evidence: Every vendor documents employer-configurable rules; the iTutorGroup case confirms employer-programmed rejection caused documented harm; employer AEDT notices describe employer-selected criteria. Limiting evidence: Some employers may rely on vendor defaults, templates, or inherited configurations they do not actively monitor. Unresolved: How many employers review or update their screening rules after initial configuration.

EXPLANATION 2: Platform-level AI architecture produces discriminatory patterns independent of employer configuration. Supporting evidence: The Mobley court found a sufficiently alleged unified policy across employers; Workday's CSM categories and HiredScore integration are common to all enabling customers; application volumes have grown far faster than job openings, amplifying the impact of any systematic bias in AI ranking. Limiting evidence: Preliminary certification is not a merits finding; employer-level configuration, population, and review practices vary materially; no validated platform-wide disparate-impact data has been publicly produced. Unresolved: The underlying audit data alleged in Mobley; whether the same Workday AI model produces materially different outcomes across employers with different configurations.

EXPLANATION 3: Market competition and applicant volume, not automation, explain most non-advancement. Supporting evidence: Applications per hire tripled from 2021 to 2024 [30]; application growth outpaced job growth by 4:1 [44]; even without AI, most applicants for any competitive posting would not receive interviews. Limiting evidence: Volume alone does not explain why some applicants report near-instantaneous rejections outside business hours [5], which is more consistent with automated disposition. Unresolved: The relative contribution of volume, qualification mismatch, and automation to non-advancement outcomes.

EXPLANATION 4: AI prioritization functions as a practical exclusion mechanism even without formal rejection. Supporting evidence: When tools assign C/D grades to ~56.5% of audited applications from applicants with known gender [14] and vendors market the ability to "skip manual screening and go directly to top candidates" [45], lower-graded applications may never receive substantive human attention. Limiting evidence: Pfizer and others state that the tool does not execute decisions and that recruiters retain full discretion [14]; no public evidence documents actual review rates by grade. Unresolved: Whether "retained discretion" translates into actual review of non-priority applications at documented scale.

THESE EXPLANATIONS ARE NOT MUTUALLY EXCLUSIVE: The most analytically defensible position is that employer configuration, platform architecture, AI prioritization, market volume, and human recruiter capacity each contribute to the observed outcomes, and that the relative contribution of each varies by employer, role, and time period.

SECTION 12

Intelligence Gaps

The following gaps represent information required to resolve the investigation's most consequential uncertainties. Each gap identifies why the information matters and what type of evidence would be needed.

GAP 1: Requisition-level outcome data by AI score or grade. This is the most important gap. No public source connects HiredScore A–D grades or Eightfold match scores with downstream outcomes (interview, advancement, rejection) at the requisition level. This data would be needed to determine whether AI prioritization functions as a practical exclusion mechanism. It exists within employer ATS systems but is not publicly available and was not produced in Mobley discovery as of the most recent court filings [6].

GAP 2: Recruiter access-event evidence. Whether a human opened and substantively reviewed an application before disposition is the core evidentiary question. ATS audit trail documentation suggests that status-change and access-event fields exist (SAP documents an application status audit trail [36]; Workday states that audit trails remain active [45]), but no public source provides actual access-event data for any private employer.

GAP 3: The Mobley-alleged 724,352-applicant audit. If this audit exists as described and contains the statistical data alleged, it would be the largest cross-employer dataset connecting ATS processing to demographic outcomes. It has not been publicly produced, independently verified, or entered as evidence [5]. Obtaining or reviewing this data would materially change the evidentiary basis of the investigation.

GAP 4: Interview-stage counts. The "interviewed" stage is the clearest indicator of substantive human consideration. EEOC MD-715 instructions require agencies to report it [27], but it is frequently missing from published agency data. No private employer examined publishes interview-stage data by posting or requisition.

GAP 5: Employer-specific threshold and queue-order documentation. At what score or grade does an employer effectively stop reviewing applications? What queue order do recruiters see? Are C/D-graded applications presented alongside A/B applications or placed in a separate, less visible queue? This information is controlled by each employer's ATS configuration and is not publicly available.

GAP 6: Post-deployment audit data for newly disclosed tools. SharkNinja disclosed that Greenhouse Talent Matching would begin NYC use on August 10, 2026, after an initial audit used test data [13]. Post-deployment audit data using historical outcomes would provide evidence about actual ranking-to-outcome relationships. This data does not yet exist.

SECTION 13

Overall Assessment

DIRECT ANSWER TO THE PRIMARY INTELLIGENCE QUESTION: The publicly documented evidence establishes that automated systems play a material and structurally significant role in determining whether job applicants receive substantive human consideration. Every major hiring platform examined provides officially documented automated screening, disposition, qualification, or prioritization capabilities, and multiple named employers deploy AI-enabled tools that classify applicants into tiers before recruiter review begins. However, the evidence also consistently establishes that employer configuration, not vendor default, is the proximate determinant of how these capabilities operate for any given requisition, and that the degree of actual human review occurring between AI prioritization and practical advancement remains the central unresolved question.

STRONGEST SUPPORTING EVIDENCE: The combination of (1) documented automated rejection, disqualification, qualification, and prioritization capabilities across the platforms examined, (2) published bias audit data showing AI grading systems classifying more than half of the audited applicant population with known gender outside priority tiers, (3) ATS aggregate data showing applications per hire exceeding 290 while interview-invitation rates hover near 3%, and (4) the iTutorGroup precedent confirming that employer-configured automation can independently produce illegal automatic rejection patterns.

WHAT REMAINS UNRESOLVED: Whether AI prioritization grades and scores translate into de facto exclusion of lower-ranked applicants. Whether any common platform architecture produces discriminatory outcomes independent of employer configuration. Whether stated human-review commitments are actually executed at scale. The content and validity of the alleged 724,352-applicant Workday audit. The relative contribution of automation versus human judgment to documented funnel contraction. Confidence level: High that documented capabilities exist and are deployed at material scale. Moderate that AI prioritization materially affects which candidates receive practical human consideration. Low that any specific quantified rejection rate attributable to platform-level AI can be established from public sources.

SECTION 14

Decision Relevance

The findings of this investigation have direct relevance to several constituencies:

For regulators and enforcement agencies: The evidence demonstrates that bias audit disclosures under NYC Local Law 144 provide useful leads but cannot serve as standalone proof of discriminatory outcomes, particularly given missing demographic data, data pooling, and the absence of threshold or outcome information [37, 43]. The federal hiring system's existing funnel structure (applied → qualified → referred → interviewed → selected) offers a model that private-sector regulators could adapt to require more granular disclosure [25, 27].

For employers deploying AI-enabled screening tools: The variation in human-review commitments across employer AEDT notices, from SharkNinja's full-application review requirement [13] to SP+/Bags' conditional and uncommitted disclosure [11], demonstrates that the specificity of review controls matters for both legal defensibility and public accountability. Employers should expect that the Mobley litigation, if it reaches the merits phase, will test whether common platform-level architecture can create employer-independent liability.

For applicants and applicant advocates: The evidence confirms that steep funnel contraction occurs before the interview stage, that AI tools classify applicants into priority and non-priority groups, and that the specific mechanisms by which applications are filtered are not publicly visible. The federal hiring system's applicant-flow data model demonstrates that transparency about qualification, referral, and selection stages is technically feasible [25, 33, 34].

For researchers and auditors: The investigation identifies specific evidence needed to advance understanding, most critically outcome data by AI score or grade at the requisition level and recruiter access-event logs. The peer-reviewed analysis of bias audit quality [37] provides a methodological framework for evaluating future audit disclosures.

For ATS vendors and AI tool developers: The investigation documents that the platform-as-gatekeeper legal theory is being tested in active litigation [8], that the EEOC has formally argued for coverage of software-mediated screening under existing anti-discrimination law [7], and that the distinction between vendor capability and employer configuration will be a central contested issue going forward.

SECTION 15

Collection Summary

The investigation's evidence base comprises 62 unique source records, each numbered in Section 17 [1]. A single source may support more than one collection-log entry when it addresses separate collection requirements or distinct evidentiary observations. Excluding the investigation's own collection log, the 61 external sources span the following categories:

Source Category	Count	Tier
Official vendor documentation (Greenhouse, Oracle, SAP, Workday, iCIMS)	18	Tier 1–2
Government records, guidance, and system documentation (OPM, EEOC, USAJOBS, NYC DCWP)	13	Tier 1
Employer-published bias audits and AEDT disclosures	7	Tier 1–2
Federal court filings and orders (Mobley v. Workday)	3	Tier 1–2
EEOC enforcement and amicus records	3	Tier 1
ATS aggregate benchmark reports (Ashby, CareerPlug)	3	Tier 2
Employer career pages and job postings	7	Tier 1
Peer-reviewed research	1	Tier 2
LinkedIn platform documentation	6	Tier 1

OPEN COLLECTION REQUIREMENTS AT INFORMATION CUTOFF: Nine collection requirements remained open as of July 27, 2026, covering recruiter-level workflow evidence, audit-log field documentation, post-deployment outcome data by AI grade, employer-specific tool activation confirmation, and underlying statistical evidence from Mobley v. Workday [1]. The most significant gap is the absence of any public source connecting AI prioritization output (scores/grades) with downstream recruiter review actions and applicant outcomes at the requisition level.

SECTION 16

Methodology

Collection Method

Sources were identified through systematic searches of vendor help centers and documentation portals, government regulatory and reporting websites, employer career pages and compliance disclosures, federal court docket filings via PACER and GovInfo, and ATS aggregate data publications. Each source was logged with a unique identifier, access date, URL, and a structured assessment of what it establishes and does not establish [1]. Collection requirements were framed as neutral intelligence questions, not propositions to confirm. The investigation did not use interviews, covert access, fabricated applications, unauthorized system testing, or private employer records.

Source Evaluation

Sources were evaluated using a tiered reliability framework:

- Tier 1 (highest): Official vendor documentation, government records, court orders, employer-published disclosures on their own platforms, and EEOC enforcement records. These sources have identifiable authorship, institutional accountability, and are published by the entity with direct knowledge of the facts described.
- Tier 2 (moderate): Vendor marketing materials, ATS aggregate benchmark reports, vendor-commissioned audits, court pleadings (allegations, not adjudicated findings), and peer-reviewed research applying independent analytical methods to secondary data.

Authority, access, specificity, transparency, recency, independence, and corroboration were assessed for each source. Where sources conflicted (e.g., Workday's assertion that it does not "recommend" applicants versus its marketing of "accelerated candidate prioritization"), both positions were recorded and the conflict was noted.

Analytical Method

Evidence was organized by collection requirement and compared across platforms, employers, and source types. Competing explanations were developed for each major observed pattern. Intelligence gaps were identified wherever the evidence was insufficient to choose between competing explanations. The analysis applied the following controls: capability was not treated as proof of use; correlation was not treated as proof of causation; allegations were not treated as findings; absence of reporting was not treated as absence of activity; and prioritization was not equated with rejection [1].

Confidence Framework

High confidence: The assessment is supported by multiple independent, authoritative sources; the evidence is consistent and specific; and alternative explanations have been considered and are less well-supported or can be ruled out.

- Moderate confidence: The assessment is supported by credible evidence from at least one authoritative source, but the evidence base is incomplete, alternative explanations remain plausible, or corroboration is partial.
- Low confidence: The assessment rests on limited, indirect, or contested evidence; significant intelligence gaps remain; or the available evidence supports multiple competing explanations with roughly equal plausibility.

SECTION 17

List of References

1. ATS and Hiring Automation OSINT Collection Log. ATS_Hiring_Automation_OSINT_Collection_Log_Scrubbed.xlsx, 2026.
2. Workday. "Workday & HiredScore's Focus on Compliance and Bias Mitigation." Workday, www.workday.com/en-us/legal/responsible-ai-and-bias-mitigation.html. Accessed 2026.
3. U.S. Equal Employment Opportunity Commission. "EEOC Sues iTutorGroup for Age Discrimination." EEOC Newsroom, 5 May 2022, www.eeoc.gov/newsroom/eeoc-sues-itutorgroup-age-discrimination.
4. U.S. Equal Employment Opportunity Commission. "iTutorGroup to Pay $365,000 to Settle EEOC Discriminatory Hiring Suit." EEOC Newsroom, 11 Sept. 2023, www.eeoc.gov/newsroom/itutorgroup-pay-365000-settle-eeoc-discriminatory-hiring-suit.
5. Mobley v. Workday, Inc. Third Amended Complaint. Case No. 3:23-cv-00770-RFL, Document 372, U.S. District Court, Northern District of California, Filed 1 July 2026.
6. Mobley v. Workday, Inc. Discovery Order. Case No. 23-cv-00770-RFL (LB), Document 340, U.S. District Court, Northern District of California, Filed 29 May 2026.
7. U.S. Equal Employment Opportunity Commission. Brief of the Equal Employment Opportunity Commission as Amicus Curiae in Support of Plaintiff. Mobley v. Workday, Inc., Case No. 3:23-cv-00770-RFL, Document 60-1, U.S. District Court, Northern District of California, Filed 9 Apr. 2024.
8. Mobley v. Workday, Inc. Order Granting Preliminary Collective Certification. Case No. 23-cv-00770-RFL, Document 128, U.S. District Court, Northern District of California, Filed 16 May 2025.
9. iCIMS. "Beyond the Hype: The Essential Need for Fair and Compliant AI Hiring." iCIMS Blog, 12 Sept. 2023, www.icims.com/blog/beyond-the-hype-the-essential-need-for-fair-and-compliant-ai-hiring/.
10. Greenhouse. "The Evolving Legal Landscape of AI in the Hiring Process." Greenhouse Blog, 20 Oct. 2023, www.greenhouse.com/blog/evolving-legal-landscape-of-ai-in-the-hiring-process.
11. SP+ / Bags. "Airline Baggage Handler Job Posting." Careers at Bags, https://careers-bags.icims.com/jobs/58467/airline-baggage-handler/job. Accessed 2026.
12. DeepHealth. "Applicant Rights and Disclosures." DeepHealth Careers, https://deephealth.com/careers/applicant-rights-disclosures/. Accessed 2026.
13. SharkNinja. "Candidate Privacy Notice New York City." SharkNinja, www.sharkninja.com/candidate-privacy-notice-new-york-city.html. Accessed 2026.
14. Pfizer. "Bias Audit Summary of the HiredScore Grade System." Pfizer, https://cdn.pfizer.com/pfizercom/2026-05/Pfizer-Bias-Audit-Summary-of-the-Hired-Score-Grade-System29MAY26.pdf, 29 May 2026.
15. Workday. "HiredScore AI for Recruiting." Workday, www.workday.com/en-us/products/talent-management/ai-recruiting.html. Accessed 2026.
16. Citizens Bank. "Digital Transformation Careers." Citizens Jobs, https://jobs.citizensbank.com/digital-transformation. Accessed 2026.
17. Citizens Bank. "Video Interview Assessment." Citizens Jobs, https://jobs.citizensbank.com/Video-Interview-Assessment. Accessed 2026.
18. McGraw Hill. "Sr Director, Applied Insights." McGraw Hill Careers, https://careers.mheducation.com/jobs/6759?lang=en-us. Accessed 2026.
19. Insmed. "Why Insmed: Our Progress, People, & How We Hire." Insmed Careers, https://careers.insmed.com/why-insmed/. Accessed 2026.
20. Insmed. "Senior Associate, Talent Discovery & Selection Operations." Insmed Careers, https://careers.insmed.com/posting/senior-associate-talent-discovery-selection-operations/R3590/. Accessed 2026.
21. Insmed. "Director, Talent & Workforce Strategy." Insmed Careers, https://careers.insmed.com/posting/director-talent-workforce-strategy/R3428/. Accessed 2026.
22. U.S. Environmental Protection Agency. "FY 2025 Annual EEO Program Status Report (MD-715)." EPA, www.epa.gov/system/files/documents/2026-06/fy-2025-epa-md-715-report-in-full.pdf, June 2026.
23. U.S. Office of Personnel Management. "Applicant Flow Data Workgroup Meeting." USA Staffing, 13 June 2019.
24. USAJOBS. "How Long Does It Take to Get a Federal Job?" USAJOBS Help Center, https://help.usajobs.gov/faq/application/how-long. Accessed 2026.
25. U.S. Office of Personnel Management. "Applicant Flow Data Warehouse: Data Relationships and Metrics." USA Staffing, 6 Sept. 2022.
26. U.S. Census Bureau. "FY 2023 Affirmative Action Plan for Recruitment, Hiring, Advancement, and Retention of Persons with Disabilities." U.S. Census Bureau, https://www2.census.gov/about/careers/working/eeo/fy2023-affirmative-action-plan.pdf, 2023.
27. U.S. Equal Employment Opportunity Commission. "Instructions to Federal Agencies for EEO MD-715." EEOC, www.eeoc.gov/federal-sector/management-directive/instructions-federal-agencies-eeo-md-715-0. Accessed 2026.

28. CareerPlug. "Recruiting Metrics & Benchmarks 2025: Industry Report." CareerPlug, www.careerplug.com/recruiting-metrics-and-kpis/, 3 Apr. 2025.
29. Ashby. "Talent Trends Report." Ashby, www.ashbyhq.com/talent-trends-report. Accessed 2026.
30. Perry, Ben. "Recruiter Productivity | 2026 Talent Trends Report." Ashby, www.ashbyhq.com/talent-trends-report/reports/2023-recruiter-productivity-trends-report, 27 Apr. 2026.
31. U.S. Office of Personnel Management. "Rule of Many Frequently Asked Questions." OPM, www.opm.gov/policy-data-oversight/hiring-information/merit-hiring-plan-resources/rule-of-many-faqs.pdf, 2025.
32. Haley, Vince, and Charles Ezell. "Merit Hiring Plan." U.S. Office of Personnel Management, www.opm.gov/chcoc/latest-memos/merit-hiring-plan.pdf, 29 May 2025.
33. U.S. Office of Personnel Management. "USA Staffing Reporting & Analytics: Reporting and Analytics Workgroup Meeting." USA Staffing, 21 June 2022.
34. U.S. Office of Personnel Management. "Reporting and Analytics Release Notes, January 7, 2026." USA Staffing, 2026.
35. U.S. Office of Personnel Management. "Access to USA Staffing Applicant Flow Data." USA Staffing, Oct. 2023.
36. SAP. "Recruiting Settings in Admin Center." SAP Help Portal, https://help.sap.com/docs/successfactors-recruiting/setting-up-and-maintaining-sap-successfactors-recruiting/recruiting-settings-in-admin-center. Accessed 2026.
37. Gerchick, Marissa Kumar, et al. "Auditing the Audits: Lessons for Algorithmic Accountability from Local Law 144's Bias Audits." The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25), ACM, 2025.
38. BABL AI Inc. "Summary of Bias Audit Results: Audit of Eightfold AI Inc.'s Eightfold Matching Model for New York City's Local Law 144." Eightfold AI, https://eightfold.ai/wp-content/uploads/eightfold-summary-of-bias-audit-results.pdf, 3 Mar. 2026.
39. DLA Piper. "TIAA New York City Local Law 144 Bias Audit Results – Workday Spotlight Tool." TIAA, 6 July 2025 (distributed 28 Jan. 2026).
40. BLDS, LLC. "Bias Audit of Morgan Stanley's Use of Eightfold Model for Scoring Applicants." Morgan Stanley, https://www.morganstanley.com/content/dam/msdotcom/en/disclosures/Bias_Audit_of_Morgan_Stanleys_Use_of_Eightfold_Model_For_Scoring_Applicants.pdf, July 2025.
41. Citizens Bank. "Annual AI Bias Audit." Citizens Jobs, https://jobs.citizensbank.com/en/Annual-AI-Bias-Audit. Accessed 2026.
42. NYC Department of Consumer and Worker Protection. "AEDT Educational Roundtable with Audit Industry Stakeholders." NYC.gov, www.nyc.gov/assets/dca/downloads/pdf/about/DCWP-AEDT-Educational-Roundtable-with-Audit-Industry.pdf, May 2023, rev. June 2023.
43. NYC Department of Consumer and Worker Protection. "Automated Employment Decision Tools: Frequently Asked Questions." NYC.gov, www.nyc.gov/assets/dca/downloads/pdf/about/DCWP-AEDT-FAQ.pdf, 29 June 2023.
44. Workday. "Workday Global Workforce Report: Job Market Tightens as AI Reshapes Hiring Processes." Workday Newsroom, https://newsroom.workday.com/2024-09-10-Workday-Global-Workforce-Report-Job-Market-Tightens-as-AI-Reshapes-Hiring-Processes, 10 Sept. 2024.
45. Workday. "AI for Office of the CHRO." Workday Webinar, https://webinars.workday.com/register/474c87c3-4454-47fa-828d-831abbfff6da/. Accessed 2026.
46. LinkedIn. "Differences Between Free and Promoted Jobs." LinkedIn Help, www.linkedin.com/help/linkedin/answer/a519737. Accessed 2026.
47. LinkedIn. "Generate the Jobs Report in Recruiter and Recruiter Lite." LinkedIn Recruiter Help, www.linkedin.com/help/recruiter/answer/a413441. Accessed 2026.
48. LinkedIn. "Application Flow for Members Applying for Your Job Posting." LinkedIn Recruiter Help, www.linkedin.com/help/recruiter/answer/a414058. Accessed 2026.
49. LinkedIn. "Reject Applicants in Recruiter." LinkedIn Marketing Solutions Help, https://www.linkedin.com/help/lms/answer/a412523. Accessed 2026.
50. LinkedIn. "Reviewing Job Applicants." LinkedIn Help, https://www.linkedin.com/help/linkedin/answer/a517574. Accessed 2026.
51. LinkedIn. "Premium Insights About Job Applicants." LinkedIn Help, https://www.linkedin.com/help/linkedin/answer/a563146. Accessed 2026.
52. Workday. "Enhance Candidate Experience with Conversational AI." Workday, https://www.workday.com/en-us/products/conversational-ai/candidate-experience.html. Accessed 2026.
53. SAP. "Start Automated Screening of Application." SAP Help Portal, https://help.sap.com/docs/successfactors-recruiting/recruiting-in-sap-successfactors-test-script/start-automated-screening-of-application-b07728388f344b82891ebeba5a786f37. Accessed 2026.
54. SAP. "Sending E-Mail Notifications to Disqualified Candidates." SAP Help Portal, https://help.sap.com/docs/successfactors-recruiting/setting-up-and-maintaining-sap-successfactors-recruiting/sending-e-mail-notifications-to-disqualified-candidates. Accessed 2026.
55. SAP. "Configuring Questions on Requisitions." SAP Help Portal, https://help.sap.com/docs/successfactors-recruiting/setting-up-and-maintaining-sap-successfactors-recruiting/configuring-questions-on-requisitions. Accessed 2026.

56. Oracle. "Create Disqualification Question in Library." Oracle Help Center, https://docs.oracle.com/en/cloud/saas/taleo-enterprise/otfru/t-createdusqualificationquestioninlibrary.html. Accessed 2026.
57. Oracle. "Guidelines for Scoring Questions." Oracle Help Center, https://docs.oracle.com/en/cloud/saas/talent-management/faimh/guidelines-for-scoring-questions.html. Accessed 2026.
58. Oracle. "Prescreening Candidates." Oracle Help Center, https://docs.oracle.com/en/cloud/saas/taleo-enterprise/20b/otfru/prescreening-candidates.html. Accessed 2026.
59. Greenhouse. "Reject a Candidate." Greenhouse Support, https://support.greenhouse.io/hc/en-us/articles/360025225732-Reject-a-candidate. Accessed 2026.
60. Greenhouse. "Auto-Advance." Greenhouse Support, https://support.greenhouse.io/hc/en-us/articles/360044105651-Auto-advance. Accessed 2026.
61. Greenhouse. "Application Rules Overview." Greenhouse Support, https://support.greenhouse.io/hc/en-us/articles/203105595-Application-rules-overview. Accessed 2026.
62. Greenhouse. "Auto-Reject." Greenhouse Support, https://support.greenhouse.io/hc/en-us/articles/360000653472-Auto-reject. Accessed 2026.

RISK INTELLIGENCE

INTELLIGENCE • EVIDENCE • IMPACT

The Collective Risk Intelligence provides analyst-led investigative intelligence, due diligence, and risk analysis for business, legal, investment, and high-stakes decision environments.

Core capabilities

- Investigative Intelligence
- Enhanced and Forensic Due Diligence
- Entity and Identity Resolution
- Supply Chain and Vendor Risk
- Litigation and Adverse Information Research
- AI Validation and Business Intelligence

Contact information

Sean Mahon
Founder & Principal

phone: 813-421-5291
email: smahon@tcriskintel.com
website: www.tcriskintel.com

Open-source demonstration report. Custom engagements are scoped according to the subject, jurisdiction, decision context, and evidentiary requirements.