2026

U.S. Algorithmic Hiring & AEDT Risk Assessment

Greenhouse ATS (Real Talent Sourcing & Match Tiering Suite)

BARKER, LEAH

1. Executive Summary

Regulatory Classification: This software is officially classified as an Automated Selection Procedure under EEOC Title VII, an Automated Employment Decision Tool (AEDT) under NYC LL 144, and a High-Impact AI System under the NIST AI RMF.

Key Findings: The system processes candidate's personal data, reviews resumes through an integrated LLM API (OpenAI sub-processor), and groups applicants into categorized match tiers (Strong, Good, Partial, Limited). While Greenhouse prevents automatic system rejections, its sorting introduces distinct U.S. legal risks regarding **automation bias** by blind trust and **proxy discrimination** from historical hiring data.

Recommendation: [X] APPROVE with conditions Final approval depends on completing an independent bias audit (Four-Fifths testing), implementing automatic data retention schedules required by state laws, and verifying OpenAI's Zero Data Retention (ZDR) agreement.

---

2. System Description

2.1 Purpose and Intended Use

The Greenhouse Applicant Tracking System manages recruitment by distributing listings to multiple job boards, gathers candidate applications, assigns applicants to tiers based on role compatibility, assists in stage notification, and uses AI to transcript and score candidate interviews.

2.2 Technical Architecture

Component	Description
Model Type	Hybrid Architecture: Custom Key Translator and OpenAI Model (GPT-4o Versions).
Training Data	OpenAI's pre-configured database, supplemented by internal history of previous candidates and successful hire records, is used for continuous refinement of the applicant list.
Input Data	Candidate resumes, cover letters, LinkedIn profile matching, and screened questionnaire responses.
Output	Categorized Match Tiers (Strong, Good, Partial, Limited), AI generated summaries of candidate qualifications, and interview grading sheets.
Update Frequency	Continuous Architecture patches from the vendor Greenhouse inc.

2.3 Deployment Context

Aspect	Description
Geographic Scope	United States (Nationwide, with specific emphasis on applicants residing in California, Colorado, Illinois, and New York City).
User Base	Internal Recruiting, Human Resources, and Interviewing Panels.
Decision Impact	High. Automatic candidate ranking acts as the primary filter determining whether a resume is reviewed by a person. If an applicant fails to match preset filters, the system automatically starts tracking their rejection record.
System Connections	Public job board websites, internal employee secure login systems, and outbound candidate email services.

---

3. Regulatory Framework Mapping

3.1 Federal and State Jurisdictional Triggers

Regulatory Body / Statute	Selected	Legal Justification
EEOC Title VII Guidance	**[X]**	Applies to any automated system or selection procedure used to screen, score, tier, or evaluate candidates that may cause systematic bias.
NYC Local Law 144 (AEDT)	**[X]**	Triggered because the Real Talent feature uses a mathematic scoring system to grade candidates and split them into filtered groups.
California Civil Rights Dept.	**[X]**	Governs automated screening software, preventing employers from shifting discriminatory blame to vendors, requiring potential employers to keep records for 4 years.
Illinois AI Video/Zip Code Acts	**[X]**	Imposes strict rules if AI analysis is used during video interviews and bans the use of ZIP codes to filter candidates by race or background.
Colorado AI Act (Consumer Protection)	**[X]**	Requires clear notifications when using hiring software and provides an option for human review if a candidate is rejected by the system.

3.2 Core Liability Determination

Under U.S. employment law, hiring tools do not have to intentionally discriminate to be illegal. If the automated ranking system creates an unfair bias against a protected group, both employer and vendor can be sued for damages.

---

4. Requirements Analysis (NIST AI RMF & State Law Alignment)

4.1 NIST Governance & Risk Management

Requirement	Current State	Gap	Remediation
Risk Identification	Privacy risk assessments provided by the vendor.	Fails to identify specific administrative bias risks.	Run test environments to observe how recruiters interact with low ranking candidates.
Tracking Performance Drift	Reviewing the vendor's official security audit report.	No data is collected to monitor the accuracy of the software's decision process over time.	Set up tracking to see where candidates drop out of the hiring process.
Safety Guards	System updates.	System is running on default rules.	Enforce a policy requiring human inspection of a randomized sample from the "Limited Match" tier.
Residual Risk Evaluation	Vendor assessment states low risk.	Fourth party data leaks are possible through vendor association with OpenAI.	Enforce strict corporate Data Processing Addendum (DPA) with the provider.

4.2 Data Governance & Proxy Auditing

Requirement	Current State	Gap	Remediation
Data Quality Criteria	The system reads standard text cleanly.	Resumes with graphics and images introduce review errors.	Ensure system alerts when a document fails to scan correctly.
Hidden Bias Checks	Vendor uses Warden AI for global bias audits.	Generic system audits cannot check for internal custom settings and controls.	Hire an independent auditor to run checks against internal system settings and controls.
Historical Hiring Data Evaluation	Preexisting data is isolated within a shared cloud infrastructure.	Historical data contains demographic skew from prior hiring practices.	Rewrite and customize skills weight for each role within the job calibration dashboard.
Data Source Tracking	Clear logs tracking input source data.	Missing third party data tracking.	Users must opt in and give permission before Chrome extensions save information from their browser.

4.3 Automated Logging & Record-Keeping

Requirement	Current State	Gap	Remediation
Automatic System Logs	Greenhouse maintains standard system audit logs.	Missing error logs from API automated alerts.	Configure external logging alerts to track API connection drops.

Decision Tracking	User modification logs are retained.	System does not log reason for overrode tier assignment.	Require a written explanation inside the dashboard when a user manually moves a candidate tier.
State Log Retention	Logs automatically purge oldest entry after standard retention timeframe.	There is no separation of logs according to region.	Create a policy to lock all system logs for 4 years to align with California Civil Rights protocols.

4.4 Transparency & Consumer Disclosure (Adverse Action / Colorado / California)

Requirement	Current State	Gap	Remediation
Candidate Disclosure	Standard privacy policies are listed.	Applicants could miss AI algorithm disclosure in complex policies.	Create a pop-up with required acceptance on the corporate website for before a candidate can apply.
Capabilities & Limits	Standard features are listed in product description.	Software does not disclose the AI margin of error to the user in real time.	Overlay a disclaimer reminding recruiters that match tiers are predictions to help assist and not absolute facts.
Adverse Action Explanations	Generic rejection templates are used.	Does not provide specific reasons for algorithm influenced rejections.	Translate system flags into clear explanations included in rejection letters.
Human Intervention Protocols	Final decisions require human interaction.	Dashboard does not provide certainty rating or alternative routes for candidates.	Establish a clear option for candidates to request a manual resume review conducted by a human.

4.5 Human Oversight Evaluation

Requirement	Current State	Gap	Remediation
Oversight Interface	Clear Recruiter Dashboard layouts.	Dashboard prioritizes top tier candidates.	Reconfigure the default dashboard sort views to randomize candidates within their tiers.
System Understanding	The dashboard is easy to navigate.	System formats are unclear to basic user.	Train the recruitment staff on how AI reads nontraditional resumes.
Override Capabilities	Users can manually move any candidate stage.	Automation bias deters regular manual overrides.	Require a periodic audit evaluating how many low tier candidates are manually advanced.
System Kill Switch	Admins can toggle off individual AI features.	Disabling features disrupts workflow.	Create a manual selection backup protocol.

4.6 Technical Robustness & Security

Requirement	Current State	Gap	Remediation
Review Accuracy	Matches human choices roughly 78% of the time.	Fails to account for AI hallucination.	Require human verification of data accuracy before candidate is approved for next stage.
Prompt Injection Defense	Active malware scanner is present.	Malware scanner treats text as passive data.	Use segmentation and least privilege to isolate the system from the internal network.
Cybersecurity / Key Rotation	SOC 2 compliant storage, data encryption at rest.	API keys are vulnerable if handled poorly.	Switch to temporary digital keys that update automatically, while verifying that all incoming data has an official cryptographic stamp from the sender.
Manipulation Controls	AI reads input as single block text.	Fails against white text prompt injection hidden in resume.	Enforce uniform styling by converting custom formatting into black and white plain text.

---

5. Civil Rights & Fundamental Fairness Impact Assessment

5.1 Protected Characteristics Assessment (U.S. Thresholds)

Characteristic	Could Be Inferred From	Risk Level	Mitigation Measure
Race/Ethnicity	Home addresses, professional association names, university locations.	High	Remove location metrics from default calibration filters.
Gender	Pronouns, single sex academic institution, gap years.	High	Enable resume blind review for early stage evaluation.
Age	Initial graduation dates, total length of career history.	High	Remove date parameters from calibration loops.
Disability	Long term employment gaps, specific accessibility request.	Med	Create a clear process where humans make the final decisions on accommodation updates.
Socioeconomic Status	Elite private university listings, unpaid internship history.	High	Standardize competency skills over specific corporate or academic names.

5.2 Vulnerable Groups Impact

Automated hiring tools that use keyword matching lists and video grading algorithms frequently filter out applicants with gaps in employment, career changers, individuals returning from extended personal or medical leave, highly skilled workers without degrees, and neurodivergent professionals.

6. Human Oversight Evaluation

6.1 Current Oversight Mechanisms

Human recruiters have complete control over changing an applicant's status. The system will never send an automatic rejection email unless a recruiter manually reviews and approves it first.

6.2 Effectiveness Assessment

Factor	Assessment	Evidence
Human reviews all outputs	Partial	Recruiters regularly review candidates in high tiers but miss low tier applicants.
Human can easily override	Yes	The user interface allows drag and drop movement across stages at any time.
Human has sufficient time	No	Large application volumes create fatigue, increasing reliance on top tier suggestions.
Human sees uncertainty	No	The system displays assigned categories directly, without showing the underlying confidence percentages.
Human trained on limits	No	Standard recruiters treat the software as an accurate search tool.
Appeals process exists	No	No direct channel exists for candidates to challenge automated tier categorizations.

6.3 Overall Assessment

- Meaningful human oversight
- **[X] Partially meaningful** (Although technical safeguards are built in, heavy workloads cause users to rely too much on automated suggestions.)
- Symbolic only

---

7. Bias & Discrimination Analysis (The Impact Ratio Audit)

7.1 Independent Bias Audit Requirements (NYC LL 144 Baseline)

The vendor or employment agency must hire an independent auditor annually to run an **Impact Ratio Test (The Four-Fifths Rule)** across gender and race classifications specifically on the automated match tiers.

7.2 The Four-Fifths Calculation Formula

To avoid legal bias violations, the tool's hiring rate for the lowest scoring group must be at least 80% of the rate for the highest scoring group. If it drops below 80%, the system cannot be used.

$$\text{Impact Ratio} = \frac{\text{selection rate for a category}}{\text{selection rate of the most selected category}}$$

OR

$$\text{Impact Ratio} = \frac{\text{scoring rate for a category}}{\text{scoring rate of the highest scoring category}}$$

---

8. Deployment Decision

8.1 Decision Matrix

Factor	Assessment	Weight	Score
Technical compliance	The vendor must have secure software systems that meet SOC 2 Type II data safety standards.	0.25	4 / 5
Civil Rights / Privacy	The system relies heavily on strict data handling agreements with third parties like OpenAI.	0.20	3 / 5
Human oversight effectiveness	High vulnerability for recruiter to rely on automated suggestions.	0.20	3 / 5
Bias Risk (Four-Fifths)	Custom settings risk accidentally tracking data that acts as a shortcut for race or gender.	0.20	2 / 5
Business benefit	Significant reduction in the time required to fill open roles and long term hiring expenses.	0.15	5 / 5
TOTAL	**Weighted Compliance Evaluation Index**	**1.00**	**3.35 / 5.00**

8.2 Recommendation

Selected Option:

- APPROVE for deployment
- **[X] APPROVE with conditions (list below)**
- DO NOT APPROVE until remediated
- PROHIBIT

8.3 Conditions for Approval (if applicable)

1. Ensure Greenhouse's agreement includes a Zero Data Retention (ZDR) policy, so candidate data is never saved or used for AI training.
2. Set up automatic schedules to erase or permanently hide candidate profiles after 4 years, meeting California Civil Rights Department requirements.
3. Add clear applicant disclosures and manual review request options to the careers website for positions located in New York, California, Colorado, and Illinois.

8.4 Monitoring Requirements

Twice a year, the compliance team must review the system to confirm security keys are being rotated, update the list of AI tools, and calculate demographic fairness ratios to ensure departments aren't accidentally discriminating.

---

9. Appendices

A. Technical Documentation Reference

- Greenhouse Software Infrastructure Security Policy and Architecture Ledger.

B. Testing Results

- Warden AI Algorithmic Fairness Public Attestation Summary Report.