

AI-Assisted Candidate Screening

How AI evaluates candidates

Prepared for DIRECTV Screening & Scheduling · September 2026

1. Purpose

The Screening & Scheduling solution uses artificial intelligence to help talent teams screen, score and shortlist candidates faster, while keeping every hiring decision firmly in human hands. This document explains, in plain terms, how our AI evaluates candidates during screening: what it looks at, how it scores, the safeguards that keep the process fair and secure, and where human judgement remains in control.

The guiding principle throughout is simple: **the AI assists, humans decide**. AI can score and recommend candidates, but it never makes the final decision to hire or reject — that action is always taken by a human recruiter.

2. The screening experience

Candidates are screened through Vera, Radancy's conversational AI screening agent. The screening conversation is structured in three stages:

- **Introduction** — Vera greets the candidate, explains the interview format and expected duration, and presents the Terms of Use for the candidate to accept before proceeding.
- **Interview** — Vera asks a sequence of questions pre-defined for the specific role. Candidates answer in the format that suits them — free text, voice note, multiple choice or video. Before moving on, Vera checks that each answer is relevant and coherent; if a response does not address the question, the candidate is asked to try again.
- **Post-interview** — Once all questions are answered, the candidate can ask their own questions, which Vera answers using the job description and a pre-defined knowledge base.

The conversation runs on a controlled, predictable workflow. Vera stays on task, redirects candidates back to the relevant question when needed, and handles every scenario through a defined set of steps — reducing ambiguity and error.

3. How the AI evaluates candidates

Radancy's scoring is designed to be structured, transparent and consistent.

Requirements and weighting. For each role, the requirements to be assessed are defined in a template, and each is given a weight according to its importance — for example Nice to have, Important or Must have. This ensures the evaluation reflects what actually matters for the role.


What the AI assesses. The score is based on all the inputs gathered for the candidate: their answers to the interview questions and any files they upload, such as a CV. For voice and video answers, the AI works from the transcript rather than the audio or images — so the assessment is driven by what the candidate said, not by subjective signals such as tone of voice or appearance.

Individual scoring with justification. Each requirement is scored individually, and every score comes with a written justification explaining the rating. This gives recruiters full transparency into how a result was reached.

No candidate-versus-candidate ranking. Candidates are never scored against one another. Each is assessed on their own merits against the role's requirements.

The final score. The overall score is a weighted calculation across all the individual requirement scores. The scoring logic can be tuned for each client and refined using feedback gathered during a testing period.

The score in practice. Each candidate receives an overall best-fit score, broken down skill by skill, with every skill accompanied by a short written note explaining its rating. This gives recruiters a clear, defensible picture at a glance:



Illustrative candidate scoring summary.

From score to action. The final score is returned to the ATS — for example, Workday — where it can drive the recruiter's business process. Clients define the thresholds that move candidates between stages, such as shortlisting stronger candidates and routing others for further review. Any automation is optional and can always be overridden by a recruiter.

Model selection and hosting. The screening model was chosen following an extensive comparison of leading large language models, selecting the strongest performer on accuracy. It runs on Microsoft Azure OpenAI, providing enterprise-grade control over data location and security.

4. Responsible AI in screening

The developer of the solution applies its established AI principles across the screening workflow, with built-in protections against bias and human ownership of every final decision.

Feature	How AI is used	Human oversight
Scorecard creation	AI suggests relevant skills and competencies for the screening questions.	The recruiter reviews, edits, customises and prioritises the skills.
AI screening agent	A conversational agent guides candidates through screening, keeps them on task and redirects to the pre-defined questions where needed. It also assesses the transcript of video answers.	The recruiter can review all candidate responses and scores and make the decision on the candidate.

Feature	How AI is used	Human oversight
Skills assessment & evaluation	AI scores candidate fit based on the CV and screening responses. Candidates are never scored against each other.	The recruiter reviews scores and decides whether to proceed, with full transparency into the AI scoring.

5. Bias mitigation

Fairness is built into the screening process through several complementary measures:

- **Anonymization & blind screening** — The Screening Tool does not use your name, photograph, age, gender, race, ethnicity, disability, ZIP code, or other protected characteristics or their proxies. It does not use video, audio, facial recognition, or other identifiers for biometric analysis.
- **Unbiased inputs** — Voice and video answers are assessed from their transcripts, and identifiers are removed before scoring, so the evaluation is not influenced by ethnicity, sex, tone or appearance.
- **Explainable AI** — Transparent, per-requirement justifications show why a candidate is recommended or not, supporting recruiter confidence and accountability.
- **Fairness monitoring** — Models are reviewed against real-world scenarios to monitor for disparate impact and selection-rate parity, so bias can be identified and corrected. Annual audits are conducted and the results be found [here](#).

6. Human oversight & decision-making

AI accelerates screening, but people stay in control:

- Recruiters review candidate responses, individual scores and the justifications behind them.
- Any automation that moves candidates based on the AI score can always be reviewed and overridden by a recruiter.
- The solution never makes the final hiring or rejection decision — that is always made by a human.
- Candidates are offered the option to opt out of AI in the hiring process.

7. Guardrails & security

The screening agent is protected against AI-specific threats and kept on a controlled path:

- **Prompt-injection protection** — System prompts and candidate inputs are strictly separated, and candidate input is never used as part of the prompt. This has been tested extensively.
- **Data-leakage protection** — The agent is engineered not to return information the candidate is not permitted to access, and sensitive information is removed before it reaches the agent's context. If a candidate asks for an API key or other sensitive detail, the agent declines.
- **Controlled workflow & answer validation** — Taira validates that each answer is relevant and coherent before accepting it, asking the candidate to retry if a response is off-topic or unclear. The conversation follows a controlled, predictable workflow, so the agent stays on its defined path and cannot be steered off task.

8. Data, storage & compliance

Data collected. From the ATS, the tool pulls the candidate's name, email and phone number, with the option to add location, CV and communication preference.

Storage. Candidate data, together with screening assessment information, is stored on AWS, and storage can be aligned to customer needs. Notifications and reminders inviting candidates to complete screening are delivered by SMS and, optionally, WhatsApp via Twilio — used for outbound messages only, with no conversations taking place on these channels.

Retention. The default data-retention period is one year and can be customised to business and regulatory requirements.

Compliance. The tool intends to be fully compliant with the EU AI Act, and all other AI and Privacy regulatory requirements, and the platform is supported by recognised security and privacy standards including SOC 2 Type II, ISO 27001 and GDPR / UK GDPR.