

greenhouse

The AI in ATS evaluation checklist



How to choose the right ATS

Every major ATS vendor now leads with AI. Feature lists are longer. Demo decks are more impressive. The pressure to adopt something is real. More AI features, however, is not a proxy for quality, and being “AI-powered” isn’t a standalone element that should make or break your procurement decision.

This checklist will serve as an evaluation supplement so you can easily assess what actually matters in an ATS.



Section 1

Start with the basics



AI works well where structure already exists.

AI should accelerate hiring processes, not fix broken or inconsistent ones. Any evaluation should consider the type of AI that's in an ATS and whether it's supporting structured hiring. Is the AI integrated in the software or is it just a surface-layer set of features?

Surface-level AI creates friction while native AI reduces it. In your walkthrough, note whether AI is an integrated part of the recruiter experience or just on the surface.

When assessing AI, consider the following three categories:

Assistive AI

Helps recruiters do tasks faster (writing, summarising, scheduling). Output is reviewed and acted on by a human.

Workflow AI

Automates steps in the hiring process (routing, triggering, reminders). Operates within defined rules.

Decision-support AI

Surfaces signals to inform hiring decisions (scoring, ranking, flagging). Directly affects who advances.

Knowing the difference not only helps with evaluation but also has ramifications when it comes to security, compliance and bias.



The basics checklist:

- Does the system enforce structured hiring by default – consistent workflows, standardised stages, role-based access and defined evaluation criteria?
- Are hiring processes consistent enough across teams and roles that AI has clean, reliable inputs to work with?
- Can hiring decisions be explained, traced and audited at each step of the process?
- Does the platform treat structured data as a first-class requirement?
- Can you clearly identify where each type of AI operates in your platform?
- Are use cases specific and tied to named product features – or described in general terms like “AI-assisted” and “intelligent recommendations”?
- When AI informs a decision, is that clearly labelled and differentiated from AI that handles a task?
- Can you show me exactly where in the workflow AI acts – and what it does?
- Does the AI operate inside the core hiring workflow – or does it require recruiters to leave the primary interface?
- Is AI functionality integrated with candidate profiles, scorecards, pipeline views and reporting – or does it live in a separate module?
- Where in the process does the AI reduce friction?
- If the AI feature were removed, would the core workflow still function?



Section 2

Ownership, explainability, governance and risk



Someone has to be accountable for every hiring decision. AI doesn't change that and should have the appropriate controls

AI in an ATS touches sensitive data, informs consequential decisions and operates across roles and geographies with varying compliance requirements. Without information on governance and compliance adherence, AI can quickly turn from an asset into a liability.

Explainability is necessary to make accountability possible. Ask a vendor to walk you through a rejected candidate's journey and explain every AI touchpoint. If they can't, your team won't be able to either. That can be risky when a candidate, regulator or an executive asks.




The basics checklist:

- Who makes the final decision at any step where AI acts?
- Can you specifically explain what the AI is doing at each stage?
- If a candidate is advanced, rejected or flagged, can that action be traced back to a human decision?
- Can AI outputs be reviewed, challenged or corrected by recruiters or TA leaders?
- Is there an audit trail that shows what the AI recommended and what the human decided?
- Are AI actions logged and auditable? Can you export or query that log?
- How is access to AI features controlled? Are permissions role-based and configurable?
- How does the platform handle PII in the context of AI? What data is used as input, where is it stored, and how long is it retained?
- How does the vendor manage AI compliance across jurisdictions with different hiring laws (EU, UK, US state-level, etc.)?
- Has the platform undergone third-party audits for bias, fairness or AI-specific compliance? Can you share the results?
- What happens to AI logs and data if you terminate the contract?



Section 3

Impact and outcomes



**“AI-powered” is a claim.
Demonstrated outcomes
are evidence.
Accountability is a must.**

AI should inform better decisions, not just faster ones. AI that bypasses structured evaluation such as scorecards, defined competencies and consistent criteria removes the accountability layer while also potentially putting the organisation at risk. Speed can't be the priority over accountability

The final evaluation layer covers three things vendors often blur together: what's live today, how their platform connects to the broader AI ecosystem you're building and what their AI means for the candidates moving through it.



The basics checklist:

- Does the AI integrate directly with structured scorecards and job-level evaluation criteria, or does it operate independently of them?
- When AI surfaces a recommendation (a score, a flag, a ranking), can the recruiter see what inputs drove it?
- Does the platform guide or make hiring decisions?
- Who's accountable for the outcome of a hiring decision?
- Does the AI improve signal quality within existing workflows, or does it bypass them?
- If a recruiter disagrees with an AI recommendation, can they override it, and is that override logged?
- For each AI capability you've been shown, ask directly: "Is this live in production, in beta or planned?"
- If a feature is on the roadmap, when was it committed? What's the expected release timeline?
- Can you speak with a current customer who is actively using the AI features you've demoed?
- Does the platform support integrations with external AI tools you already use or plan to use?
- Are there open APIs, MCPs or established integration pathways for AI-related data (candidate signals, scoring outputs, workflow triggers)?
- How does the vendor manage conflicts or overlaps between their native AI and third-party AI tools?
- How does the platform disclose AI use to candidates? Is that disclosure configurable?
- Has the vendor tested their AI for disparate impact across demographic groups? Are those results available?



Section 4

Common mistakes to avoid



Most evaluations go wrong before the RFP is issued:

Evaluating by feature volume.

More AI features does not equate with better AI. A platform with five well-integrated AI capabilities is more useful than a platform with 20 loosely connected ones.

Treating all AI the same.

Assistive AI, workflow AI and decision-support AI carry different risks and require different governance. Evaluate them separately.

Prioritising speed over decision quality.

AI that processes candidates faster but makes them less visible – or less fairly evaluated – is not an improvement.

Accepting vague “AI-powered” claims.

Every claim should be tied to a specific feature, a specific workflow step and a specific outcome. If a vendor can't do that, the feature isn't ready.

Assuming “AI-first” means better outcomes.

AI-first architecture is a design choice, not a quality signal. Ask what problem the AI is solving – not how prominent it is in the product.

Ignoring governance and auditability. Don't get blinded by speed and flashy marketing – having accountability and governance is a must.

Not asking about candidates.

The hiring process is experienced by candidates, not just recruiters. AI that damages candidate trust or creates unexplained friction has real downstream costs.



AI is changing how hiring decisions get made. The question is whether that change makes decisions better – or just faster and harder to explain.

The right system does three things. It supports decisions that are consistent across roles, teams and geographies. It keeps humans accountable at every step where it matters. And it builds the kind of trust – with candidates, hiring managers and regulators – that holds up when it's tested.

That's the standard.
Evaluate to it.

greenhouse

Greenhouse is the hiring platform built on structured hiring. Everything we build – including AI – is designed to make hiring decisions more consistent, explainable and defensible. [Learn more about how Greenhouse approaches AI.](#)