



AI tools Due Diligence Guide

Protecting your business*

AI is reshaping recruitment. CV screening, candidate matching, predictive attrition, interview scheduling — vendors are pitching solutions that promise to cut costs and accelerate time-to-fill. But for every tool that delivers, many do not. And in recruitment, the cost of getting it wrong extends beyond wasted budget to regulatory exposure, reputational damage, and real harm to the candidates you are supposed to serve. However, there are some phenomenal tools out there that will help drive growth & efficiency, but selecting the right partner to invest in demands a complex procurement and decision-making process.

The purpose of this guide is to arm you with the questions that will help support you through that process, with emphasis on protecting your business. The legal and regulatory risk in AI procurement sits primarily with you as the buyer, not the vendor and contracts will not protect you if the underlying tool creates discriminatory outcomes under anti-discrimination legislation or breaches data legislation as outlined in Australia's Privacy Act or GDPR for example.

This isn't set and forget. New regulations and guidelines regarding the ethical use of AI are constantly emerging, so applying this level of rigour to your procurement process will enable you to monitor ongoing compliance of your tools and technologies with emerging law and standards.

Before we dive in though, some thoughts to align on to ensure you're set up for success:

Problem definition	Is there a clearly defined baseline (current process or recruiter performance) against which improvement will be measured – what will good look like?
Data	Is your current ATS data clean enough e.g. are key fields (e.g. job titles, skills, locations) standardised rather than free-text? Implementing AI on bad data is just lipstick on a pig. Are outcomes (e.g. placement, rejection, progression) clearly defined and reliably recorded, are those outcomes relayed to the vendor and used to update the training data for the tool? Does your historical data reflect potential bias or skewed decision-making patterns?
ROI	Is there a full cost model including implementation, legal, and change management? What's your exit strategy if it doesn't produce results? Bear in mind the impact on engagement of the teams running the pilot if it fails.
Proof of concept	Has the PoC been run on your own data with a pre-set pass threshold? Only looking at vendor supplied data is going to give a skewed view.
Pilot	Did the pilot hit the target KPI? Is adoption/change plan in place? Implementation and change planning plays a critical role in ensuring success.
Scale	Are governance, drift monitoring, and vendor review cadence all defined? This won't be set and forget – there will need to be ongoing management.

**This document is intended as a practical operational guide and does not constitute legal advice. For current and specific guidance on compliance for your organisation consult a qualified legal adviser – see final page for details.*

1

Transparency

Can you and your teams explain what's happening through the process?

Can the vendor explain in plain language how the AI scores/ranks candidates? Specifically, what personal data is processed and how it influences candidate scores or rankings? If not, how will your team explain it to candidates?

Does the tool operationally support human-review requests, or is the tool used to make a decision (including doing anything which is substantially and directly related to making a decision) that may significantly affect an individual's rights? If this is the case, this information must be disclosed in your Privacy Policy (AUS) as of December 2026. There may also be requirements to disclose this information under other laws, e.g. GDPR.

Are the system / AI-generated decisions logged and auditable? This is currently only a requirement under GDPR, but tracking this retrospectively could be a real and challenge so it might be prudent to build this in now.

Key notes

- Satisfy yourself that the vendor is open (glass box vs black box!) If the vendor is a "black box" or you are otherwise not comfortable to explain solely (or mostly) automated decisions, it is important to ensure that there is meaningful human involvement in the decision
-

2

Bias & Ethics

AI models reflect the data they were trained on. In recruitment, that data has historically encoded discrimination.

What training data was used? UK / US / AU or mixed? US training data encodes US labour market patterns: different job title taxonomies, skills frameworks, and compensation norms. A model trained predominantly on US data will likely underperform on Australian candidates (think “server” – very different meanings) — sometimes in ways that correlate with protected characteristics.

Has bias been tested across all protected characteristics not just gender, ethnicity and age, but others such as disability, age, and religion?

Have you run disparity analysis on your own candidate data vs just the vendor demo?

Does the model update continuously to stay current with market changes? (otherwise, it risks becoming outdated - new job titles for example)

Does the model learn from your hiring decisions over time? If so, how is reinforcement of existing bias prevented?

Key notes

- AI uses a massive amount of data to train itself – were there biases inherent when it did so?
 - Are you satisfied that the vendor mitigated (and continues to mitigate) that risk?
 - How are you ensuring (contractually) that the vendor continues to provide you with the latest and greatest version of their model?
-

3

Data & Compliance

This is the most complex area and the one most frequently under-interrogated during procurement. Compliance risk is multi-layered and, critically, it sits with you.

Where is data processed and stored by the vendor? Who has access to that data and where are they based?

Is any data transferred outside Australia? Do ALL of the entities handling that data also comply with Australian Privacy requirements?

Does the tool scrape data from social media or job networking sites? If so, how is this compliant with data legislation (and your contract with that platform!) in your region?

Does candidate data feed into the vendor training model? If so, how is it anonymised?

What is the data deletion policy on contract exit? Has that been written into your agreement?

Are data retention policies (both internally and within vendor systems) clearly documented and defined? Is candidate data only retained for as long as necessary, in line with applicable legal requirements?

Will the vendor support you in completing Privacy Impact Assessment (AUS)?

Key notes

- It's your responsibility to know what happens to your data
 - The vendor doesn't know what your contractual arrangements are with social media platforms
-

4

Candidate obligations

Your obligation to inform and protect candidates begins before the AI touches a single CV.

What candidate-facing privacy notice language does the vendor supply if relevant?

Is there an operational process for candidates to request access to AI-generated scoring data and challenge or contest AI assisted decisions?

Are candidates informed at point of data collection that AI will be used? This should be stated in your process / Privacy Policy

5

Vendor fit

Check what you're buying really works....

Have you tested on your own live data — not just on vendor sample sets and are vendor claims (e.g. accuracy, bias reduction) independently validated, or based solely on vendor-provided evidence?

Does it outperform your current process or human consultant benchmark? This will be your success metric baseline.

What are the SLAs, remedies for breach or "service unavailable", and exit / data offboarding terms?

Key points

- Candidate obligations are one of the areas most likely to change over time – so future proofing is important here
 - Vendor demos are designed to impress so you really need to get under the hood
-



Key takeaways:

You own the risk – not the vendor

If you can't explain a decision, you can't defend it

Your data drives the bias, not the demo data

Your responsibility extends beyond your environment and includes your vendor's data use

If it's not delivering better results than your consultants delivered, pull it!

For more information:

apac@thesatoripartnership.com

[Check out our Data Edge Framework](#)

For further legal advice contact:

Tanvi.mehtakrensel@squirepb.com

[Tanvi Mehta Krensel, Partner Squire Patton Boggs](#)