

**SPRING
2026**



**Is Your AI Strategy
Audit-Ready?
Turning Compliance
into a Competitive
Advantage**

VONQ

VIEW

AI IS NO LONGER AN EXPERIMENT IN HIRING.

IT IS INFRASTRUCTURE.

Across Talent Acquisition, AI is now embedded in screening, scoring, interviewing, and decision support. It has helped organisations manage scale, reduce manual workload, and improve efficiency. But as adoption has accelerated, so has scrutiny.

Regulators have caught up. In the European Union, AI used in hiring is formally classified as high-risk. In the United States, independent bias audits are now mandatory in jurisdictions like New York City and Colorado. Across the UK and other global markets, new rules are reshaping what responsible AI deployment looks like in practice.

This marks a structural shift. The question is no longer whether to use AI in hiring. The question is whether your AI can withstand audit, legal challenge, and public scrutiny. Many systems cannot.

Early-generation hiring AI was designed to filter candidates quickly, often operating as opaque “black boxes” with limited explainability and weak governance. That model is becoming a liability. Organisations are now accountable not only for hiring outcomes, but for how those outcomes are produced.

COMPLIANCE IS NO LONGER A LEGAL AFTERTHOUGHT.

IT IS A DESIGN REQUIREMENT.

Leading organisations are responding by adopting a higher standard - building AI systems that are transparent, human-controlled, independently audited, and continuously monitored. This shift is not just about reducing risk. It is about creating hiring infrastructure that can be trusted, defended, and scaled globally.

This edition of The VONQ View explains what audit-ready AI actually looks like, how the global regulatory landscape is evolving, and the practical steps organisations can take to protect themselves.

It also outlines VONQ’s approach, including the HEAT framework and our EQO AI agents, which were designed from the outset to meet the highest compliance standards through independent assurance and continuous oversight.

Because in the next phase of AI adoption, success will not be defined by automation alone.

It will be defined by trust and by whether your AI can be explained, audited, and defended when it matters most.

From innovation to control: AI maturity is bringing scrutiny

It's no exaggeration to say that Artificial Intelligence (AI) has today become omnipresent in so many of our recruitment and talent acquisition (TA) workflows, processes including CV screening, chatbots, automated assessments and ranking tools. According to the [World Economic Forum¹](#), close to 90% of employers are now using AI of some form in their hiring processes.

But the narrative around AI in hiring has changed from “how do we make the most of it?” to “how do we control it?” Initial industry enthusiasm turned to anguish as one of the defining examples of the so-called “dangers” of AI generally centred on a perceived bias in hiring practices. This anguish rapidly converted into governance as jurisdiction after jurisdiction enacted strict regulations designed to rein in those risks.

Especially for organisations operating in multiple industries and geographies, compliance with this regulation is as exacting as it is essential.

Non-compliance with this global regulatory environment can wreak significant damage across any organisation. It corrodes trust, erodes employer brand and reputation and can lead to expensive fines and punishing penalties.

So how can you stay on the right side of this ever shifting line?



90%

...of employers are now using AI of some form in their hiring processes.



The Spring edition of the VONQ View Quarterly unpacks the myriad of compliance demands on organisations of all sizes to help recruiters and TAs navigate this emerging regulatory landscape.



Importantly, it outlines VONQ's “HEAT Framework” of principles and disciplines that sets organisations up for compliance success wherever you are operating, and however that landscape shifts.

What is AI's 'black box screening' problem?

The primary concern with using AI in the recruitment process is what's known as "black box screening". This describes the opacity of early AI systems and the perceived mysteriousness of how they made candidate recommendations - either to shortlist or to screen-out - and whether that decision was fair or not.

Not only should the conclusions reached by any AI system on a candidate be explainable and transparent, but it's equally essential that human involvement is self-evident.

When it comes to AI in TA, the biggest risk is bias, of which there are four kinds to guard against:

01

Data bias: poor representation of diversity in terms of age, gender or ethnic heritage

02

Algorithmic bias: Flaws in the code or logic that inadvertently amplify historical prejudices and systemic inequalities

03

User bias: developer prejudice that bakes in discriminatory bias into the data models

04

Deployment bias: Mismatches between the environment where a model was built and the real-world context where it is used.

We explore bias in AI-based recruiting in more detail in our **June 2025 edition of the VONQ View:**

https://www.VONQ.com/wp-content/uploads/2025/06/VONQVIEW_june.pdf

Navigating the global regulatory landscape

Wherever you are in the world, and depending on whether AI is playing an internal or external, candidate-facing role, you will be accountable to a complex matrix of compliance demands designed to keep AI under control. Importantly, these laws apply where the candidate is based, not where the employer is hiring from.

In setting regulatory frameworks, governments around the world face a difficult balance between protecting consumers and citizens while simultaneously enabling and fostering innovation and competition.

This landscape is ever shifting as governments change political hue or economic priority. Postures will pivot as what we know and understand about AI evolves and as new threats and risks emerge.



Designing for AI Compliance: The VONQ Approach

To illustrate what compliant AI product design looks like in practice, we unpack our own approach as a practical example - reflecting our decision to build compliance-readiness into EQO from the ground up.

We have ensured that AI is only responsible for relieving human TAs of much of the administrative hinterland that comes before a decision - screening, scoring, assessment and even interviewing. It does not decide. Humans still decide - and in that sense AI adds value without taking away control.

Watching the burgeoning regulation movement developing around the world, and seeing the role that auditing would clearly play in governance in so many countries, we made a decision to be very proactive in making EQO totally audit-ready.

„The moment AI makes decisions instead of supporting them, then we know a line has been crossed.“



-  **Independent assurance**
-  **Continuous bias auditing**
-  **Post-market monitoring**
-  **Transparent reporting**

A useful rule of thumb here is:

YOU CAN'T DEFEND WHAT YOU DO NOT UNDERSTAND.

To ensure EQO operates fairly in real-world use, VONQ engaged Warden AI, an independent AI audit and assurance provider. Their role is to independently evaluate whether EQO behaves equitably across protected characteristics using external datasets and established bias detection methodologies.

Warden does not endorse or influence the system under test. Warden conducts ongoing monthly audits, publishing results through a transparent public dashboard which demonstrates how EQO performs against using multiple complementary techniques designed to detect both structural and emerging bias, and how it would measure up to the most stringent AI GRC standards.



Disparate Impact Analysis – assessing whether outcomes differ across demographic groups



Counterfactual Analysis – testing whether changing protected attributes alters results



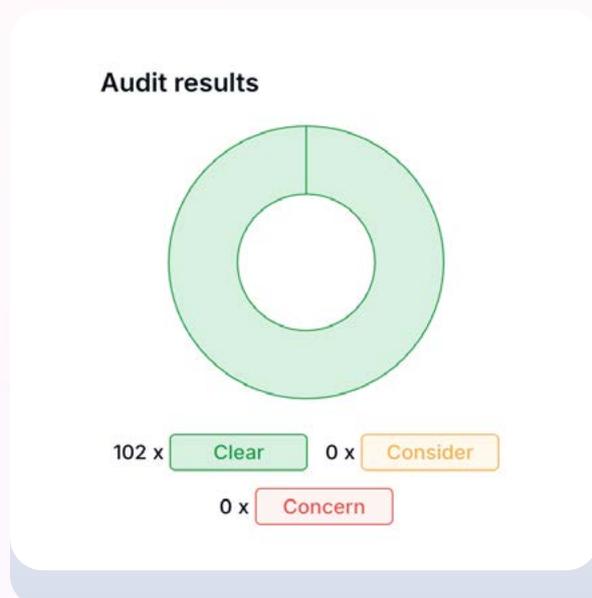
Independent external datasets representing diverse candidate populations



Continuous monitoring to detect model drift and unexpected bias

Importantly, this is not a static certification. Continuous monitoring ensures fairness remains stable as the system evolves, a critical requirement for trustworthy AI in operational hiring environments.

The future of hiring AI will not be defined by automation alone, but by verifiable trust.



Bias & Fairness	
Sex bias	Clear
Race/Ethnicity bias	Clear
Intersectional bias (Sex X Race/Ethnicity)	Clear
Age bias	Clear
Disability bias	Clear
Religion bias	Clear
Sexual orientation bias	Clear
Veteran status bias	Clear
English proficiency bias	Clear
Pregnancy status bias	Clear
National origin bias	Clear
Criminal history bias	Clear
Medical conditions bias	Clear
Gender identity bias	Clear
Marital status bias	Clear

“Skills-first” versus matching

One aspect of AI-based hiring that has perpetuated biased outcomes in early iterations was ‘matching’, an historically common approach to ranking candidates by virtue of how well they align with or mirror the existing workforce. What’s worse is that these types of models are quite opaque when it comes to revealing how they reach ranking decisions. Matching models replicate your past workforce. Skills-first models help you build your future workforce.

That’s why at VONQ, we take a skills-first approach. Instead of mirroring the existing workforce, our AI Agents apply a structured point-based assessment across 15 defined job-relevant vectors - evaluating candidates against objective indicators of capability rather than background similarity.

This enables potential to be recognized where traditional matching would likely screen candidates out, expanding access to qualified talent that conventional models may systematically exclude.

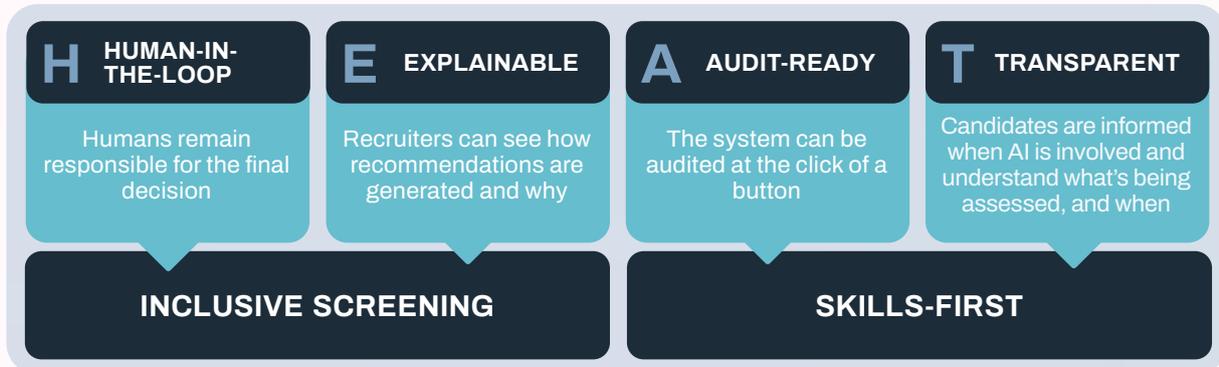
By focusing on demonstrable capability, organisations can identify individuals who may not share the same experience on paper but clearly possess the competencies required to succeed in the role - recognising transferable strengths such as structured communication, conflict resolution and process discipline that are directly relevant to customer-facing, administrative or operational positions.

Because candidate evaluations are grounded in clearly defined, job-related vectors, organisations can articulate how and why recommendations are made - rather than relying on opaque similarity matching.

Not only is it fairer, more transparent and easier to explain, it is more effective at making better hires.

A 2025 LinkedIn survey found that companies leaning into skills-based searches are 12% more likely to make a quality hire.

The HEAT Framework: What compliant AI looks like



Regardless of geography, there are common elements that often arise that are consistent across all legal requirements that can be considered a governance “high water mark”.

To meet this standard, VONQ developed “The HEAT Framework”, a set of four core principles that together establish a “safety-by-design” approach to the use of AI in the hiring process designed to protect candidates and applicants from AI prejudice:

- 01 Human-in-the-loop:** Humans remain responsible for the final decision.
- 02 Explainable:** Recruiters can see how recommendations are generated and why.
- 03 Audit-ready:** The system can be audited at the click of a button.
- 04 Transparent:** Candidates are informed when AI is involved and understand what's being assessed, and when.

One fundamental design principle that matters more than most is: screening people in, not out. Making these systems inclusive rather than exclusive means that by default all candidates are given equal access to opportunity.

Passing the compliant AI “trust test”

Here are some questions you and your team should be able to answer about your AI hiring process without hesitation:

- Are candidates clearly informed when AI is used?
- Can we explain the logic of the system in plain English?
- Can recruiters see and audit how recommendations are made?
- Is there independent evidence of fairness and bias testing?
- Do humans always make the final decision?

If any of those feel uncomfortable, that's a signal the compliance-worthiness of your “ADM” might require further scrutiny. If those questions continue to cause discomfort, it is worth thinking about inviting independent consultants to take a look. It is certain that the cost of such a project will still be less than any fines you might incur if you ignore the problem.

One approach can be to proactively publish an AI framework to which you can be held publicly accountable. In 2022 Microsoft published its “Responsible AI Standards” which includes conducting regular risk assessments, implementing data protection measures and prioritizing transparency and accountability in decision-making. Google published its own “AI Principles” in 2023 which are very similar, stressing the importance of fairness, transparency and privacy in AI development.

View-from-the-CTO's Office: AI Bias and how to take it on

“At VONQ, we believe technology should elevate human potential and broaden access to opportunity. As AI becomes part of how organizations identify, screen, and interview talent, mitigating bias and ensuring fairness is essential.

That is why independent validation from Warden confirming our alignment with fairness standards across regulated jurisdictions is so important. It shows that rigorous governance, transparency, and continuous oversight are central to building technology that HR leaders can trust.

The recent class-action lawsuit filed against Eightfold AI under the Fair Credit Reporting Act and related laws highlights the legal and ethical risks when AI hiring tools are used without clear transparency and candidate protections.

This regulatory moment makes it even more important for HR leaders to scrutinize how AI systems are deployed and governed. For HR leaders, deploying AI responsibly means anchoring these technologies in accountability and explainability and respecting the rights of candidates. Fairness and compliance are not just about meeting regulatory requirements. They are about ensuring candidates experience the process as transparent and just.

Mitigating bias and governing AI hiring responsibly requires ongoing effort and clear commitments. Our approach includes:



Regular independent auditing of models and outcomes to detect disparities



Data practices that are representative and inclusive to reduce the risk of reinforcing historical bias



Structured human oversight so that AI augments informed judgement rather than replaces it

AI can help HR teams streamline processes and uncover qualified talent more effectively, but fairness, transparency, and candidate trust are indispensable.

Our commitment is to build with thoughtfulness, validate with rigor, and operate with transparency so that HR leaders can adopt AI with confidence, build more diverse teams, and ensure every candidate is evaluated fairly.”



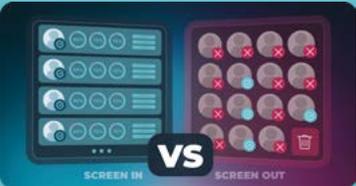
Conclusion: Building for the future

The use of AI in recruitment is here to stay, and is transforming how we hire now and into the future. That change is manifestly positive - for both candidates and TAs. But there are risks, and governments are right to regulate its use to protect citizens. Systems that aren't compliant are not sustainable. Non-compliance is a dangerous risk for trust, employer brand and for hiring outcomes. Non-compliance can also hit the bottom line...hard.

The standards emerging out of today's global GRC initiatives are responsible and ethical and organisations should embrace the opportunity to align with those standards to both protect employees - current and future - and to guarantee better hiring outcomes. In the coming years, the defining question will not be "Does your AI work?" It will be: "Can you defend it?"



What to read next:



Black Box Hiring Is Becoming A Liability

by our CEO, Ritu Mohanka



Screening In: The Only Hiring Model You Can Defend

written by Özge Özgülyüz

1. "Hiring with AI doesn't have to be so inhumane", World Economic Forum: <https://www.weforum.org/stories/2025/03/ai-hiring-human-touch-recruitment/>
2. "Shaping Europe's Future: The EU AI Act", The European Commission: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
3. "Virginia Governor Vetoes Artificial Intelligence Bill HB 2094: What the Veto Means for Businesses", Ogletree Deakins: <https://ogletree.com/insights-resources/blog-posts/virginia-governor-vetoes-artificial-intelligence-bill-hb-2094-what-the-veto-means-for-businesses/>
4. "2024 Generative AI Consumer Trust Survey", KPMG: <https://kpmg.com/us/en/media/news/generative-ai-consumer-trust-survey.html>

Key Takeaways



As concerns rose about the prevalence of bias in the use of AI in hiring, governments around the world have moved to introduce governance, regulation and compliance (GRC)



Those concerns mainly centre on the concept of “black box screening” which describes the use of opaque systems where decisions on candidates are not transparent and cannot be easily explained



A best-practice approach to such a varied compliance landscape should target the “high water mark” of regulation in order to ensure consistent compliance, and should include:

- “Human-in-the-loop” oversight
- Explainability
- Audit-readiness
- Transparency (HEAT)



Our approach to AI systems - such as our EQO team of AI agents introduced late last year - can be a helpful model to follow and includes:

- Screening in, not out
- Skills first vs matching
- Independent assurance (Warden AI)
- Continuous bias auditing
- Post-market monitoring
- Transparent reporting.



A skills-first approach bakes in fairer and more transparent outcomes than a traditional matching approach

Start Your AI Advantage Today- with EQO

The teams that win this year are already transforming their hiring now. Discover how EQO AI Agents can accelerate your workflows, screen and shortlist candidates in hours instead of days, and give you the competitive edge before the new year begins.

